

MAINVIEW® SRM EasyPOOL User Guide and Reference

Version 7.1

March 28, 2002



Copyright © 1994–2002 BMC Software, Inc., as an unpublished work. All rights reserved.

BMC Software, the BMC Software logos, and all other BMC Software product or service names are registered trademarks or trademarks of BMC Software, Inc. IBM and DB2 are registered trademarks of International Business Machines Corp. All other registered trademarks or trademarks belong to their respective companies.

THE USE AND CONTENTS OF THIS DOCUMENTATION ARE GOVERNED BY THE SOFTWARE LICENSE AGREEMENT ENCLOSED AT THE BACK OF THIS DOCUMENTATION.

Restricted Rights Legend

U.S. GOVERNMENT RESTRICTED RIGHTS. UNPUBLISHED–RIGHTS RESERVED UNDER THE COPYRIGHT LAWS OF THE UNITED STATES. Use, duplication, or disclosure by the U.S. Government is subject to restrictions set forth in FAR Section 52.227-14 Alt. III (g)(3), FAR Section 52.227-19, DFARS 252.227-7014 (b), or DFARS 227.7202, as amended from time to time. Send any contract notices to Contractor/Manufacturer:

BMC Software, Inc.
2101 CityWest Blvd.
Houston TX 77042-2827
USA

Contacting BMC Software

You can access the BMC Software Web site at <http://www.bmc.com>. From this Web site, you can obtain general information about the company, its products, special events, and career opportunities. For a complete list of all BMC Software offices and locations, go to <http://www.bmc.com/contact.html>.

USA and Canada

Address	BMC Software, Inc. 2101 CityWest Blvd. Houston TX 77042-2827
Telephone	713 918 8800 or 800 841 2031
Fax	713 918 8000

Outside USA and Canada

Telephone	(01) 713 918 8800
Fax	(01) 713 918 8000

Customer Support

You can obtain technical support by using the Support page on the BMC Software Web site or by contacting Customer Support by telephone or e-mail. To expedite your inquiry, please see “Before Contacting BMC Software,” below.

Support Web Site

You can obtain technical support from BMC Software 24 hours a day, seven days a week by accessing the technical support Web site at <http://www.bmc.com/support.html>. From this site, you can

- read overviews about support services and programs that BMC Software offers
- find the most current information about BMC Software products
- search a database for problems similar to yours and possible solutions
- order or download product documentation
- report a problem or ask a question
- subscribe to receive e-mail notices when new product versions are released
- find worldwide BMC Software support center locations and contact information, including e-mail addresses, fax numbers, and telephone numbers

Support via Telephone or E-mail

In the USA and Canada, if you need technical support and do not have access to the Web, call 800 537 1813. Outside the USA and Canada, please contact your local support center for assistance. To find telephone and e-mail contact information for the BMC Software support center that services your location, refer to the Contact Customer Support section of the Support page on the BMC Software Web site at www.bmc.com/support.html.

Before Contacting BMC Software

Before you contact BMC Software, have the following information available so that a technical support analyst can begin working on your problem immediately:

- product information
 - product name
 - product version (release number)
 - license number and password (trial or permanent)
- operating-system and environment information
 - machine type
 - operating system type, version, and service pack or program temporary fix (PTF)
 - system hardware configuration
 - serial numbers
 - related software (database, application, and communication) including type, version, and service pack or PTF
- sequence of events leading to the problem
- commands and options that you used
- messages received (and the time and date that you received them)
 - product error messages
 - messages from the operating system, such as `file system full`
 - messages from related software

Contents

About This Book	xi
-----------------------	----

Chapter 1	What is EasyPOOL?
-----------	-------------------

Chapter 2	How to Use EasyPOOL
-----------	---------------------

Chapter 3	EasyPOOL Functions
-----------	--------------------

DASDPOOL—Allocate Data Sets to DASD Volume Pools	3-3
Usage Notes	3-17
DSNCHECK—Enforce Data Set Naming Standards	3-31
Usage Notes	3-34
FDRASIST—Pool Data Sets Recalled With FDR/ABR	3-39
Usage Notes	3-43
FORCECAT—Force Catalog Disposition for New Data Sets	3-45
Usage Notes	3-46
MODDELET—Remove Specific Volume for MOD-DELETE Data Set Deletion	3-48
Usage Notes	3-49
OPENEMPT—Initialize Empty Data Sets	3-50
Usage Notes	3-51
OPTBLKSZ—Set Optimum Blocksize	3-54
Usage Notes	3-56
SETEXPDT—Set Expiration Date for Data Set	3-59
Usage Notes	3-62
SMSACSTE—Test ACS Routines	3-66
Usage Notes	3-67
SMSMCREN—Identify Data Set Rename Requiring Management Class Change	3-70
Usage Notes	3-71
SMSSELECT—Control Selection of Volumes from SMS Storage Groups	3-74
Usage Notes	3-85
SPACCONV—Change Allocation Units to Blocks	3-94
Usage Notes	3-96

SPACLIMI—Limit the Size of Space Allocation Requests	3-99
Usage Notes	3-101
SPACRLSE—Release Data Set Space at Close	3-104
Usage Notes	3-105
SPACSQTY—Set Primary and Secondary Space for Data Set.	3-108
Usage Notes	3-112
STKSUPP—Split Unit Affinity Specifications Over Devices.	3-116
Usage Notes	3-118
SUPJSCAT—Suppress JOBCAT and STEPCAT DD Statements	3-122
Usage Notes	3-123
SUPVOLRF—Suppress Volume References	3-125
Usage Notes	3-126
TAPECOMP—Control Tape Data Set Compaction (IDRC).	3-129
Usage Notes	3-130
TAPEDEFR—Defer Tape Mounts	3-132
Usage Notes	3-133
TAPEPOOL—Assign Tape Data Sets to Pools.	3-134
Usage Notes	3-137
VIOALLOC—Control VIO Allocation.	3-140
Usage Notes	3-141
VSAMCNTL—Enforce VSAM Allocation Standards	3-143
Usage Notes	3-147

Chapter 4

How to Define Critical Data Set Lists

How to Use the CritList Option.	4-3
---	-----

Chapter 5

How to Define DFSMS Subpools

How to Use the SMS Pools Option	5-5
---	-----

Glossary

Index

List of Figures

Figure 3-1	MVOLPER View	3-26
Figure 3-2	Possible System Storage Configuration	3-28
Figure 3-3	MVOLPER View	3-89
Figure 5-1	Possible System Storage Configuration for SMS-Managed Subpools	5-3

List of Tables

Table 2-1	Application Groupings of EasyPOOL Functions	2-2
Table 2-2	EasyPOOL Functions and System Parameters Cross-Reference	2-3
Table 3-1	System Parameters—DASDPOOL	3-4
Table 3-2	VSAM Component Separation Conditions	3-19
Table 3-3	System Parameters—DSNCHECK	3-31
Table 3-4	System Parameters—FORCECAT	3-45
Table 3-5	System Parameters—OPTBLKSZ	3-54
Table 3-6	System Parameters—SETEXPDT	3-59
Table 3-7	ACS Variables Displayed by SMSACSTE	3-68
Table 3-8	System Parameters—SMSSELCT	3-75
Table 3-9	Volume Selection Techniques	3-86
Table 3-10	System Parameters—SPACCONV	3-94
Table 3-11	System Parameters—SPACLIMI	3-99
Table 3-12	System Parameters—STKSUPP	3-116
Table 3-13	System parameters—TAPEPOOL	3-134
Table 5-1	SET Statement SMS-Managed Subpool Parameter Quick Reference .	5-2
Table 5-2	INC/EXC Statement Pool Parameter Quick Reference	5-2

About This Book

This book contains detailed information about MAINVIEW® Storage Resource Manager EasyPOOL by BMC Software and is intended for storage administrators.

To use this book, you should be familiar with the following:

- OS/390 operating system
- job control language (JCL)
- Interactive System Productivity Facility (ISPF)

How This Book Is Organized

This book is organized as follows. In addition, an index and glossary appear at the end of the book.

Chapter/Appendix	Description
Chapter 1, "What is EasyPOOL?"	provides a brief overview of the EasyPOOL product
Chapter 2, "How to Use EasyPOOL"	lists the functions grouped according to general applications
Chapter 3, "EasyPOOL Functions"	provides detailed information and examples of each EasyPOOL function
Chapter 4, "How to Define Critical Data Set Lists"	provides detailed information about how to define critical data sets
Chapter 5, "How to Define DFSMS Subpools"	provides detailed information about how to define DFSMS subpools

Related Documentation

BMC Software products are supported by several types of documentation:

- online and printed books
- online Help
- release notes and other notices

In addition to this book and the online Help, you can find useful information in the publications listed in the following table. As “Online and Printed Books” on page xii explains, these publications are available on request from BMC Software.

Category	Document	Description
MAINVIEW common documents	<i>OS/390 and z/OS Installer Guide</i> <i>MAINVIEW Installation Requirements Guide</i> <i>MAINVIEW Common Customization Guide</i> <i>Using MAINVIEW</i> <i>MAINVIEW Administration Guide</i> <i>Implementing Security for MAINVIEW</i>	provide instructions for installing, configuring, using, and administering MAINVIEW
MAINVIEW SRM customization documents	<i>MAINVIEW SRM Customization Guide</i>	provides instructions for configuring and customizing MAINVIEW SRM for OS/390 including EasyPOOL
core documents	<i>MAINVIEW SRM User Guide and Reference</i>	provides information common to all MAINVIEW SRM products and high-level navigation
	<i>MAINVIEW SRM Reference Summary</i>	provides a reference of global parameters, filter list and rule list parameters, and functions
messages	<i>MAINVIEW SRM Messages</i>	provides hardcopy of messages that are also available online
supplemental documents	release notes, flashes, technical bulletins	provides additional information about the product

Online and Printed Books

The books that accompany BMC Software products are available in online format and printed format. If you are a Windows or Unix user, you can view online books with Acrobat Reader from Adobe Systems. The reader is provided at no cost, as explained in “To Access Online Books.” You can also obtain additional printed books from BMC Software, as explained in “To Request Additional Printed Books.”

To Access Online Books

Online books are formatted as Portable Document Format (PDF) files. You can view them, print them, or copy them to your computer by using Acrobat Reader 3.0 or later. You can access online books from the documentation compact disc (CD) that accompanies your product or from the World Wide Web.

In some cases, installation of Acrobat Reader and downloading the online books is an optional part of the product-installation process. For information about downloading the free reader from the Web, go to the Adobe Systems site at <http://www.adobe.com>.

To view any online book that BMC Software offers, visit the support page of the BMC Software Web site at <http://www.bmc.com/support.html>. Select a product to access the related documentation.

To Request Additional Printed Books

BMC Software provides printed books with your product order. To request additional books, go to <http://www.bmc.com/support.html>.

Release Notes and Other Notices

Printed release notes accompany each BMC Software product. Release notes provide current information such as

- updates to the installation instructions
- last-minute product information

In addition, BMC Software sometimes provides updated product information between releases (in the form of a flash or a technical bulletin, for example). The latest versions of the release notes and other notices are available on the Web at <http://www.bmc.com/support.html>.

Conventions

This section provides examples of the conventions used in this book and explains how to read ISPF panel-flow diagrams and syntax statements.

General Conventions

This book uses the following general conventions:

Item	Example
information that you are instructed to type	Type SEARCH DB in the designated field.
specific (standard) keyboard key names	Press Enter .
field names, text on a panel	Type the appropriate entry in the Command field.
directories, file names, Web addresses	The BMC Software home page is at www.bmc.com .
nonspecific key names, option names	Use the HELP function key. KEEPDICTIONARY option
MVS calls, commands, control statements, keywords, parameters, reserved words	Use the SEARCH command to find a particular object. The product generates the SQL TABLE statement next.
code examples, syntax statements, system messages, screen text	//STEPLIB DD The table <i>table_name</i> is not available.
emphasized words, new terms, variables	The instructions that you give to the software are called <i>commands</i> . In this message, the variable <i>file_name</i> represents the file that caused the error.
single-step procedures	» To enable incremental backups, type y and press Enter at the next prompt.

This book uses the following types of special text:

Note: Notes contain important information that you should consider.

Warning! Warnings alert you to situations that could cause problems, such as loss of data, if you do not follow instructions carefully.

Tip: Tips contain useful information that may improve product performance or that may make procedures easier to follow.

Syntax Statements

Syntax statements appear in Courier. The following example shows a sample syntax statement:

```
COMMAND KEYWORD1 [KEYWORD2|KEYWORD3] KEYWORD4={YES|NO}  
      file_name...
```

The following table explains conventions for syntax statements and provides examples:

Item	Example
Items in italic type represent variables that you must replace with a name or value. Use an underscore for variables with more than one word.	<i>dtsbackup control_directory</i>
Brackets indicate a group of options. You can choose at least one of the items in the group, but none of them is required. Do not type the brackets when you enter the option. A comma means that you can choose one or more of the listed options. You must use a comma to separate the options if you choose more than one option.	[<i>table_name, column_name, field</i>]
Braces enclose a list of required items. You must enter at least one of the items. Do not type the braces when you enter the item.	{ <i>DBD_name table_name</i> }
A vertical bar means that you can choose only one of the listed items. In the example, you would choose either <i>commit</i> or <i>cancel</i> .	{ <i>commit cancel</i> }
An ellipsis indicates that you can repeat the previous item or items as many times as necessary.	<i>column_name . . .</i>

Chapter 1 What is EasyPOOL?

EasyPOOL is one of the products that make up the MAINVIEW SRM line of storage management products. EasyPOOL functions provide control over data set allocation and enforcement of allocation and naming standards. EasyPOOL functions operate at the operating system level to intercept normal job processing, thus providing services without any JCL changes.

The following are some of the benefits and uses of EasyPOOL:

- allocates new data sets to pools instead of individual volumes
- allocates temporary data sets to VIO
- limits allocated space
- enforces standards by controlling:
 - data set naming
 - cataloging
 - data set expiration date usage
 - space allocation units (blocks versus tracks/cylinders)
 - release of unused space
 - tape data set compaction
- optimizes block sizes
- suppresses JOBCAT and STEPCAT DD statements
- organizes tape devices into pools
- redirects tape data sets to DASD and vice versa
- performs affinity separation of tape volumes in/out of silos
- isolates critical data sets

In addition, volume selection can be based on user-defined criteria (maximum space, EMC volume, performance, and so on). Data sets can be pooled based on historical performance data collected for up to two weeks.

Chapter 2 How to Use EasyPOOL

EasyPOOL consists of a set of functions that help you

- pool data sets
- enforce data center standards
- enhance tape functionality
- enhance DFSMS functionality
- enhance data set allocation efficiency
- enhance volume selection
- isolate critical data sets
- run DASD Performance Optimization (DPO)
- run Historical DASD Performance Optimization

Using EasyPOOL involves activating and controlling EasyPOOL functions. Refer to the *MAINVIEW SRM User Guide and Reference* for details.

Note: Pools must be defined prior to setting up EasyPOOL functions. Refer to the *MAINVIEW SRM User Guide and Reference* for information about defining pools.

The following table lists the functions grouped according to general applications. See “EasyPOOL Functions and System Parameters Cross-Reference” on page 2-3 for a matrix of EasyPOOL functions and the system parameters that support them.

Table 2-1 Application Groupings of EasyPOOL Functions

If you need to...	Consider using...
pool data sets	<p>"DASDPOOL—Allocate Data Sets to DASD Volume Pools" on page 3-3</p> <p>"FDRASIST—Pool Data Sets Recalled With FDR/ABR" on page 3-39</p> <p>"MODDELET—Remove Specific Volume for MOD-DELETE Data Set Deletion" on page 3-48</p> <p>"SUPVOLRF—Suppress Volume References" on page 3-125</p> <p>"VIOALLOC—Control VIO Allocation" on page 3-140</p>
enforce data center standards	<p>"DSNCHECK—Enforce Data Set Naming Standards" on page 3-31</p> <p>"FORCECAT—Force Catalog Disposition for New Data Sets" on page 3-45</p> <p>"SUPJSCAT—Suppress JOBCAT and STEPCAT DD Statements" on page 3-122</p> <p>"VSAMCNTL—Enforce VSAM Allocation Standards" on page 3-143</p>
enhance tape functionality	<p>"STKSUPP—Split Unit Affinity Specifications Over Devices" on page 3-116</p> <p>"TAPECOMP—Control Tape Data Set Compaction (IDRC)" on page 3-129</p> <p>"TAPEDEFER—Defer Tape Mounts" on page 3-132</p> <p>"TAPEPOOL—Assign Tape Data Sets to Pools" on page 3-134</p>
enhance DFSMS functionality	<p>"SMSACSTE—Test ACS Routines" on page 3-66</p> <p>"SMSMCREN—Identify Data Set Rename Requiring Management Class Change" on page 3-70</p> <p>"SMSSELECT—Control Selection of Volumes from SMS Storage Groups" on page 3-74</p>
enhance data set allocation efficiency	<p>"OPENEMPT—Initialize Empty Data Sets" on page 3-50</p> <p>"OPTBLKSZ—Set Optimum Blocksize" on page 3-54</p> <p>"OPTBLKSZ—Set Optimum Blocksize" on page 3-54</p> <p>"SETEXPDT—Set Expiration Date for Data Set" on page 3-59</p> <p>"SPACCONV—Change Allocation Units to Blocks" on page 3-94</p> <p>"SPACLIMI—Limit the Size of Space Allocation Requests" on page 3-99</p> <p>"SPACRLSE—Release Data Set Space at Close" on page 3-104</p> <p>"SPACSQTY—Set Primary and Secondary Space for Data Set" on page 3-108</p>
enhance volume selection	<p>"DASDPOOL—Allocate Data Sets to DASD Volume Pools" on page 3-3</p> <p>"SMSSELECT—Control Selection of Volumes from SMS Storage Groups" on page 3-74</p>
isolate critical data sets	<p>"DASDPOOL—Allocate Data Sets to DASD Volume Pools" on page 3-3</p> <p>"SMSSELECT—Control Selection of Volumes from SMS Storage Groups" on page 3-74</p>
run DASD Performance Optimization (DPO)	<p>"DASDPOOL—Allocate Data Sets to DASD Volume Pools" on page 3-3</p> <p>"SMSSELECT—Control Selection of Volumes from SMS Storage Groups" on page 3-74</p>
run Historical DASD Performance Optimization (HISTDPO)	<p>"DASDPOOL—Allocate Data Sets to DASD Volume Pools" on page 3-3</p> <p>"SMSSELECT—Control Selection of Volumes from SMS Storage Groups" on page 3-74</p>

Some functions within EasyPOOL are affected by specific system parameters. The following table lists system parameters and the EasyPOOL functions affected by them.

Table 2-2 EasyPOOL Functions and System Parameters Cross-Reference

System Parameter	DASDPOOL	DSNCHECK	FDRASIST	FORCECAT	MODDELET	OPENEMPT	OPTBLKSZ	SETEXPDT	SMACSTE	SMSMCREN	SMSSELCT	SPACCONV	SPACLIMI	SPACRLSE	SPACQTY	STKSUPP	SUPJSCAT	SUPVOLRF	TAPECOMP	TAPEDEFR	TAPEPOOL	VIOALLOC	VSAMCNTL
BLKINPUT							X																
BLKOLDSR							X																
CHECK	X																						
CRITLIST	X										X												
DASDGENR	X	X		X				X															
DMYUNIT	X																				X		
DP_RENAME	X																						
HISTDAYS	X										X												
JCLEXT	X																						
JCLUREQ	X																						
PROCOLD	X																						
REJECT	X	X											X								X		
SMS_ALLOC											X												
SMS_EXTEND											X												
STKSCR																	X						
TAPEGENR		X		X				X													X		
TRKCYL							X					X											
TRKLEN							X					X											

Chapter 3 EasyPOOL Functions

MAINVIEW SRM storage management services are divided into functions. SMFUNCxx activates EasyPOOL functions and controls message and tracing activity. Functions provide all the runtime services of EasyPOOL. SMFUNCxx points to members SMFLSTxx and SMRLSTxx, which select resources and control the operation of the functions.

Functions are defined in SET statements. You can change parameters

- by editing the member directly
- by using the EasyPOOL function SET commands
- through the Functions option on the Parmlib pop-up menu of the EZSRM Menu (see Chapter 6, “How to Define and Activate Functions” in the *MAINVIEW SRM User Guide and Reference*).

Listed below are the functions available in EasyPOOL. They are described in greater detail in the remaining pages of this book.

DASDPOOL—Allocate Data Sets to DASD Volume Pools	3-3
DSNCHECK—Enforce Data Set Naming Standards	3-31
FDRASIST—Pool Data Sets Recalled With FDR/ABR	3-39
FORCECAT—Force Catalog Disposition for New Data Sets	3-45
MODDELET—Remove Specific Volume for MOD-DELETE Data Set Deletion	3-48
OPENEMPT—Initialize Empty Data Sets	3-50
OPTBLKSZ—Set Optimum Blocksize	3-54
SETEXPDT—Set Expiration Date for Data Set	3-59
SMSACSTE—Test ACS Routines	3-66
SMSMCREN—Identify Data Set Rename Requiring Management Class Change	3-70
SMSSELECT—Control Selection of Volumes from SMS Storage Groups	3-74
SPACCONV—Change Allocation Units to Blocks	3-94
SPACLIMI—Limit the Size of Space Allocation Requests	3-99
SPACRLSE—Release Data Set Space at Close	3-104
SPACSQTY—Set Primary and Secondary Space for Data Set	3-108

STKSUPP—Split Unit Affinity Specifications Over Devices.	3-116
SUPJSCAT—Suppress JOBCAT and STEPCAT DD Statements	3-122
SUPVOLRF—Suppress Volume References	3-125
TAPECOMP—Control Tape Data Set Compaction (IDRC).	3-129
TAPEDEFR—Defer Tape Mounts	3-132
TAPEPOOL—Assign Tape Data Sets to Pools.	3-134
VIOALLOC—Control VIO Allocation.	3-140
VSAMCNTL—Enforce VSAM Allocation Standards.	3-143

EasyPOOL functions do not operate for DFDSS jobs.

DASDPOOL—Allocate Data Sets to DASD Volume Pools

The DASDPOOL function is used to assign data sets to pools of DASD volumes. Pools allow multiple volumes to be managed as single entities. Pools allow data sets to be grouped and managed by their characteristics. Pool storage groups are a key component of DFSMS.

Considerations in DASD pooling include:

- Use of storage versus private volumes
- Volume allocation for non-specific requests
- Acceptance/rejection of requests for allocation on specific volumes
- Grouping of data sets by usage, by size, or by physical characteristics

Assignment of data sets to pools should be centrally controllable by the data center, without requiring JCL changes or any type of source changes. The storage administrator should be able to easily manage the storage of the entire organization.

DASDPOOL allows you to set standards for data sets' volume residency. A collection of data sets, selected by filter and rule list parameters from a variety of job and data set properties, can be directed to reside on the volumes in a single pool or in multiple pools for more flexibility in space utilization and management.

Critical Data Set Isolation

EasyPOOL allows you to identify a list of critical data sets in your environment that should never be placed on the same volume as any other data set in the list. EasyPOOL may also use data gathered by StorageGUARD about the EMC Symmetrix environments to ensure that critical data sets are not only on separate logical volumes but are placed on separate physical volumes to reduce data set contention.

Note: StorageGUARD must be installed and running to access EMC Symmetrix data.

It is important to consider, however, that CRITDSN is resource intensive and should be used only for a *small* list of critical data sets. It should not be used without considering the system impact. See Chapter 4, “How to Define Critical Data Set Lists,” for information on defining the SMCRITxx definition parameter.

Note: The DASDPOOL function must be specified in the SMFUNCxx member to be available and must be set active to provide services.

System Parameters

The system parameters that affect DASDPOOL are described in Table 3-1.

Table 3-1 System Parameters—DASDPOOL (Part 1 of 2)

Parameter	Purpose
CHECK	Specifies candidate volume replacement during allocation for multi-volume data sets. If multiple volumes are requested for a data set, VSAM or non-VSAM, the CHECK parameter controls whether the first volume or all volumes are immediately assigned from the assigned pool.
CRITLIST	Specifies the suffix of an SMCRITxx parameter member.
DASDGENR	Specifies the only generic DASD names that are to be processed for specific functions if no space requirements are specified. Standard pooling functionality allows the JCL UNIT parameter to specify a pool name. If this is not desired, DASDGENR can be used to intercept data set allocations without space information. Such allocations will only be considered for subsequent processing if the unit is found in an internal table (3380, 3350, and so on) or if the unit is found in the DASDGENR list. DASDGENR affects functions DASDPOOL, DSNCHECK, SETEXPDT, and FORCECAT. If no DASDGENR table is coded, all DASD allocations will be processed.
DMYUNIT	Defines the conversion of a nonexistent UNIT parameter to a valid UNIT parameter only if JCLEXT=YES.
DP_RENAME	Indicates whether DASDPOOL is to be processed on renamed data sets.
HISTDAYS	Specifies the number of days of data to retrieve from the StorageGUARD historical performance data collector and temporarily store in a data space. Values are 0-14 (7 or 14 recommended), which indicate the number of days of data to be retrieved. The default value is 0, which indicates that the data space is created, but no historical performance data is gathered and no HISTDPO pooling can be performed. If this value is changed, the historical performance data collector (SVSGP) must be stopped and restarted. For a valid HISTDPO test, do not specify a number of HISTDAYS higher than the total number of days collected in the historical performance data collector.
JCLEXT	Specifies if EasyPOOL will obtain volume and unit information after the operating system accesses the catalog. This is useful when PROCOLD=YES and/or DISP=MOD is specified.
JCLUREQ	When PROCOLD=YES is specified and EasyPOOL is analyzing a DD statement for an existing data set, the JCLUREQ parameter determines if UNIT information is required. If JCLUREQ=YES is specified, MAINVIEW SRM intercepts the DD statement only when the JCL specifies UNIT parameter. This allows the data center to correct JCL that uses an invalid unit parameter for existing data sets without analyzing DD statements that are correctly specified. EasyPOOL examines every DD statement associated with an existing data set if you specify JCLUREQ=NO. Note: JCLUREQ=NO needs to be specified when processing UNIT=AFF groups since the unit field in the SIOT consists of blanks.

Table 3-1 System Parameters—DASDPOOL (Part 2 of 2)

Parameter	Purpose
PROCOLD	Specifies if EasyPOOL will intercept DD statements that specify OLD allocations. EasyPOOL always intercepts NEW and MOD allocations. EasyPOOL also intercepts DD statements that specify the UNIT parameter when PROCOLD=YES is specified, which is useful for correcting questionable JCL. For example, assume DASDPOOL pools a data set to a TEST volume. If a later step wants to use the data set and specifies UNIT=PROD, the job receives a JCL error. You can correct this situation by specifying PROCOLD=YES.
REJECT	Controls termination of processing at the first or last data set rejected by REJECT=YES in DASDPOOL or DSNCHECK. If termination is to take place on the first rejected data set, code REJECT=FIRST. If all data sets are to be processed before control is returned to OS/390 allocation, code REJECT=LAST.

Rule List Parameter Quick Reference

Parameter	Description
CANDIDATE=YES/NO	Candidate volumes accepted (VSAM data sets)
CRITEMC=YES/NO	Specifies whether the volume meeting CRITDSN criteria includes EMC physical volumes
CRITBIAS=nnnn	Defines the number of data sets that can reside on the volume before the current allocation
CRITFAIL=YES/NO	Defines allocation if a volume meeting the criteria cannot be found
CRITLIST=xxxxxxx	Name of the table containing allocation volumes
DPORDEF=nnn	Default response time for device selection
DPORMAX=nn	Maximum response time target for device selection
DPORMIN=nn	Minimum response time target for device selection
DPORSEP=nn	Separation factor for device selection
DPOWIND=nnnn	Interval for device selection based on performance
EVENTID=xxxxx	Identifies a user event defined in an SMEVNTxx member
HDPORDEF=nnn	Default response time for device selection
HDPORMAX=nn	Maximum response time target for device selection
HDPORMIN=nn	Minimum response time target for device selection
HDPORSEP=nn	Separation factor for device selection
HDPOSTIM=nnnn	Starting time each day for a range of snapshots to be analyzed
HDPOETIM=nnnn	Ending time each day for a range of snapshots to be analyzed
HDPODAYS=(MO,TU,WE,TH,FR,SA,SU)	Day(s) of the week used in determining performance statistics for pooling
NVOLMAX=YES/NO	Limits volumes to the number in the assigned pool

Parameter	Description
NVOL= <i>nn</i>	Number of volumes and units that can be allocated to a data set
NVOLINDX= <i>nn</i>	Number of volumes that can be allocated to a VSAM index component
POOL=(<i>xxxxxxxx,....</i>)	1-15 pool names (from SMPPOOLxx) (1-8 characters)
RAIDDEVTYPE= <i>xxxxxxxx</i>	Allows the user to require a specific RAID device type for an allocation (EMC/RDFEMC/MIRROREMC/PARITYEMC) (requires StorageGUARD) For RAIDDEVTYPE to be honored, a VOLSEL routine <i>must be</i> assigned.
REJECT= YES/NO	Rejects recommended volume assignment
SEP= YES/NO/ASIS	Specifies volume separation of data/index VSAM components
SUPVOL= YES/NO	Suppresses request for allocation to a specific volume
UNIT= <i>xxxxxxxx</i>	Unit name (esoteric or generic) 1-8 characters
USEVOL=STOR/PRIV/ALL	Directs volume allocation to a specific type (STOR, PRIV, all)
VFORCE= YES/NO	Adds standard component suffixes (DATA, INDEX) to non-JCL defined VSAM file names
VOLSEL= <i>xxxxxxxx</i>	Volume selection criteria (BESTFIT, CRITDSN, DPO, HISTDPO, MAXSPACE, PERCENT)

Rule List Parameter Explanations

CANDIDATE=

Purpose: Specifies whether the IDCAMS DEFINE VOL=() list for a VSAM data set is used as an indicator of secondary volume allocation. If CANDIDATE=YES, the *number* of DEFINED volumes is used to select that many candidate volumes from the MAINVIEW SRM pool. If CANDIDATE=NO, no secondary volumes are assigned or available through normal OS/390 processing (but SPACVOLA can be used to assign secondary volumes from a MAINVIEW SRM pool). Note that with specification of the CANDIDATE parameter, the candidate volumes are assigned from the pool, not from those volumes specified in the IDCAMS DEFINE VOL list. CANDIDATE overrides NVOL and NVOLINDX. The CANDIDATE parameter does not apply to JCL-defined VSAM data sets.

Syntax: CANDIDATE=YES/NO

Default: CANDIDATE=NO

CRITBIAS=

Purpose: Defines the number of data sets from the assigned CRITLIST that can reside on the volume before the current allocation.

Syntax: CRITBIAS=*xxxx*
where *xxxx* is a number 0–9999.

Default: None

Note: This parameter works only in conjunction with VOLSEL=CRITDSN.

CRITEMC=

Purpose: Specifies whether the volume meeting CRITDSN criteria includes EMC physical volumes. If yes, data sets specified with CRITDSN should not reside on the same EMC Physical volume nor the same MVS Logical volume. If no, data sets specified with CRITDSN may reside on the same EMC physical volume. Default is no.

Syntax: CRITEMC=*YES/NO*

Default: CRITEMC=NO

Note: This parameter works only in conjunction with VOLSEL=CRITDSN.

CRITFAIL=

Purpose: Defines the allocation process if a volume meeting the CRITDSN criteria cannot be found. If Yes, allocation fails; if No, allocation is made to the best available volume. Like USELIM, if an acceptable volume cannot be found in the first pool, the best volume from the first pool is saved, and the system tries the next pool(s). If no acceptable volume is found, CRITFAIL is processed on the saved volume.

Syntax: CRITFAIL=*YES/NO*

Default: CRITFAIL=NO

Note: This parameter works only in conjunction with VOLSEL=CRITDSN.

CRITLIST=

Purpose: The table that contains a list of critical data sets.

Syntax: CRITLIST=xxxxxxxx

Default: None

Note: This parameter works only in conjunction with VOLSEL=CRITDSN.

DPORDEF=

Purpose: Specifies the default response time for device selection based on DASD performance. The default value is used when response time information has not been accumulated for a device, such as when it has just been brought online.

DPORDEF is useful only when specified in conjunction with DPORMIN/MAX.

Syntax: DPORDEF=*nnn*

where *nnn* specifies the default response time in milliseconds that will be substituted for unavailable information for a specific device.

Default: DPORDEF=50

DPORMAX=

Purpose: Specifies the maximum response time objective for device selection based on DASD performance.

DPORMAX is normally specified in conjunction with DPORMIN.

Syntax: DPORMAX=*nn*

where *nn* specifies the maximum response time in milliseconds that will be considered for device selection based on performance.

Default: None

DPORMIN=

Purpose: Specifies the minimum response time objective for device selection based on DASD performance.

DPORMIN is normally specified in conjunction with DPORMAX.

Syntax: DPORMIN=*nn*

where *nn* specifies the minimum response time in milliseconds that will be considered for device selection based on performance.

Default: None

DPORSEP=

Purpose: Specifies the response time value that is used to force selection of different volumes for data sets in the same jobstep. When multiple data sets are allocated in a single jobstep and device selection based on DASD performance is specified for some or all, this separation factor is added to the current response time for previously used volumes in order to increase the likelihood that new allocations will go to a different volume.

Syntax: DPORSEP=*nn*

where *nn* specifies the response time in milliseconds that will be used to enhance data set separation across volumes.

Default: DPORSEP=10

DPOWIND=

Purpose: Specifies the window (length of the performance interval) that is used to analyze the response characteristics of DASD devices for selection based on performance.

Syntax: DPOWIND=*nnnn*

where *nnnn* specifies the number of seconds over which the performance analysis will be made.

Default: DPOWIND=1800 (this is the maximum data retained in the performance tables)

EVENTID=

Purpose: Specifies the identifier assigned to a user event in SMEVNTxx. This parameter will cause an event to be generated from this function.

Note: If EVENTID= is used on an FLST SET statement with MODE=INACT, the event will still be issued.

Syntax: EVENTID=xxxxx

where xxxxx is the 5-character string specified on the EVNTID parameter in SMEVNTxx.

Default: None

HDPORDEF=

Purpose: Specifies the default response time for device selection based on DASD performance. The default value is used when response time information has not been accumulated for a device, such as when it has just been brought online.

HDPORDEF is useful only when specified in conjunction with HDPORMIN/MAX. This parameter is available only when VOLSEL=HISTDPO.

Syntax: HDPORDEF=nnn

where nn specifies the default response time in milliseconds that will be substituted for unavailable information for a specific device.

Default: HDPORDEF=50

HDPORMAX=

Purpose: Specifies the maximum response time objective for device selection based on DASD performance.

HDPORMAX is normally specified in conjunction with HDPORMIN. This parameter is available only when VOLSEL=HISTDPO.

Syntax: HDPORMAX=*nn*

where *nn* specifies the maximum response time in milliseconds that will be considered for device selection based on performance.

Default: None

HDPORMIN=

Purpose: Specifies the minimum response time objective for device selection based on DASD performance.

HDPORMIN is normally specified in conjunction with HDPORMAX. This parameter is available only when VOLSEL=HISTDPO.

Syntax: HDPORMIN=*nn*

where *nn* specifies the minimum response time in milliseconds that will be considered for device selection based on performance.

Default: None

HDPORSEP=

Purpose: Specifies the response time value that is used to force selection of different volumes for data sets in the same jobstep. When multiple data sets are allocated in a single jobstep and device selection based on DASD performance is specified for some or all, this separation factor is added to the current response time for previously used volumes in order to increase the likelihood that new allocations will go to a different volume. This parameter is available only when VOLSEL=HISTDPO.

Syntax: HDPORSEP=*nn*

where *nn* specifies the response time in milliseconds that will be used to enhance data set separation across volumes.

Default: HDPORSEP=10

HDPOSTIM=

Purpose: Specifies the starting time (each day) for a range of snapshots to be analyzed. This allows analysis of data from the same time range, or shift, each day. This parameter is available only when VOLSEL=HISTDPO.

Syntax: HDPOSTIM=*nnnn*

where *nnnn* specifies a time in 24-hour format.

Default: HDPOSTIM=0000

HDPOETIM=

Purpose: Specifies the ending time (each day) for a range of snapshots to be analyzed. This allows analysis of data from the same time range, or shift, each day. This parameter is available only when VOLSEL=HISTDPO.

Syntax: HDPOETIM=*nnnn*

where *nnnn* specifies a time in 24-hour format.

Default: HDPOETIM=2359

HDPODAYS=

Purpose: Specifies the day(s) of the week to be used in determining performance statistics for pooling. This parameter is available only when VOLSEL=HISTDPO.

Syntax: HDPODAYS=(*MO, TU, WE, TH, FR, SA, SU*)

Default: HDPODAYS=(*MO, TU, WE, TH, FR, SA, SU*)

NVOL=

Purpose: Specifies the number of volumes and units that can be allocated to a data set.

Syntax: NVOL=*nn*

where *nn* is a number in the range of 1–59 for non-VSAM data sets and 1–20 for VSAM data components.

Default: None

NVOLINDEX=

Purpose: Specifies the number of volumes that can be allocated to a VSAM index component. Note that CANDIDATE overrides NVOLINDEX for non-JCL defined VSAM data sets.

Syntax: NVOLINDEX=*nn*

where *nn* is a number in the range of 1–20.

Default: None

NVOLMAX=

Purpose: Specifies whether the number of volumes that can be assigned to a data set is limited to the number of volumes in the pool that the data set resides in. Note that volume counts in alternate pools are not included in NVOLMAX.

Syntax: NVOLMAX=*YES/NO*

Default: NVOLMAX=NO

POOL=

Purpose: Specifies the name of the pool to receive the data set. If multiple pools are specified, the first pool with sufficient space for the data set is assigned. (The first pool specified is the primary pool; all others are alternate pools.) Note that a pool specification is not required in certain circumstances, such as when rejecting an allocation (REJECT=YES) or allowing a specific volume request (SUPVOL=NO).

Syntax: POOL=xxxxxxx or POOL=(xxxxxxx,xxxxxxx,...)

where xxxxxx is a pool name defined in SMPOOLxx. If the poolname is not defined, the refresh or start-up will fail. Up to 15 pools can be assigned within parentheses.

Default: None

RAIDDEVTYPE=

Purpose: Allows the user to require a specific RAID device type for an allocation. If a particular hardware type is chosen, the candidate list of volumes will be limited to devices that have the requested attribute flag enabled. This parameter requires StorageGUARD to be installed and running.

Note: For RAIDDEVTYPE to be honored, a VOLSEL routine *must be* assigned.

Syntax: RAIDDEVTYPE=xxxxxxx

where xxxxxx is one of the following:

EMC The device must be an EMC RAID device.

RDFEMC The device must be EMC and have the RDF flag enabled.

MIRROREMC The device must be EMC and have the MIRROR flag enabled.

PARITYEMC The device must be EMC and have a parity flag enabled.

Default: None

REJECT=

Purpose: Specifies that the recommended device assignment (from OS/390 or DFSMS) is to be rejected. No DASDPOOL processing is performed; the request is terminated immediately, and message SVM3311I is issued. Note that this causes the allocation (and the job) to fail.

Syntax: REJECT=*YES/NO*

Default: REJECT=NO

SEP=

Purpose: Specifies whether the data and index components of a non-JCL defined VSAM key-sequenced data set are allocated to separate volumes in a pool.

CANDIDATE=YES must also be specified for VSAM component separation.

Syntax: SEP=*YES/NO/ASIS*

Default: None

SUPVOL=

Purpose: Specifies whether an allocation request for a specific volume is suppressed. If SUPVOL=YES, the request for a specific volume is ignored, and the data set is assigned to a pool. If SUPVOL=NO, the specific volume request is honored and the data set is not pooled. This parameter applies to both VSAM and non-VSAM data sets.

Syntax: SUPVOL=*YES/NO*

Default: SUPVOL=YES

UNIT=

Purpose: Specifies a unit name (esoteric or generic) to be used by the data set during allocation.

Syntax: UNIT=*xxxxxxxx*

where *xxxxxxxx* is a valid unit name (esoteric or generic) 1–8 characters long.

Default: None

USEVOL=

Purpose: Specifies the type of volume (storage, private, or either) that satisfies non-specific allocation requests. USEVOL=STOR directs the data set allocation to a storage-mounted volume; USEVOL=PRIV directs the data set allocation to a private volume. (There is currently no specification for PUBLIC-only volumes.) USEVOL=ALL directs the allocation to the first available volume, regardless of use attribute.

Syntax: USEVOL=*STOR/PRIV/ALL*

STOR storage mounted volume
PRIV privately mounted volume
ALL volume of any mount type

Default: USEVOL=ALL

VFORCE=

Purpose: Specifies naming conventions for non-JCL defined VSAM components by adding standard component suffixes (DATA, INDEX) to VSAM data set cluster names.

Syntax: VFORCE=*YES/NO*

YES overrides unspecified or invalid VSAM component names by appending .DATA and .INDEX qualifiers to the cluster name

NO does not modify VSAM component names

Default: VFORCE=NO

VOLSEL=

Purpose: Specifies that volume selection within a pool should be based on current volume space availability.

Syntax: *VOLSEL=BESTFIT/CRITDSN/DPO/HISTDPO/MAXSPACE/PERCENT*

Specifies that the volume with the smallest contiguous extent that satisfies the primary allocation should be selected.

BESTFIT	Specifies that the volume with the smallest contiguous extent that satisfies the primary allocation should be selected.
CRITDSN	Controls allocations by specifying data sets that should not reside on the same volume.
Warning!	CRITDSN is resource intensive and should be used only for a <i>small</i> list of critical data sets. It should not be used without considering the system impact.
DPO	Specifies volume selection based on performance statistics accumulated by MAINVIEW SRM.
HISTDPO	Specifies volume selection based on historical volume performance.
Note:	HISTDPO requires more resources than normal allocations.
MAXSPACE	Specifies that the volume with the largest single contiguous extent should be selected.
PERCENT	Specifies that the volume with the largest amount of free space should be selected.

Default: None

Parameters Not Supported

The following filter list parameters are not supported for DASDPOOL:

- GDGVER
- VALUE

Usage Notes

DASDPOOL provides services for non SMS-managed data sets with a status disposition of NEW or MOD. (The function SMSSELECT performs this service for DFSMS-managed data sets.)

Data sets can be assigned to a single pool, or to a collection of up to 15 pools. When multiple pools are specified, DASDPOOL checks each pool in the order specified to locate sufficient space for the data set. All pools specified must be defined in SMPOOLxx or the refresh or start-up will fail.

DASDPOOL allows pool names to be used as OS/390 esoteric device names, without definition as such in OS/390. Data sets assigned to the DASDPOOL pool name specified in the JCL UNIT parameter will be assigned to a volume in that pool. For example, if pool MAXSIZE is defined in SMPOOLxx, then:

```
//NEWTRANS DD DSN=data-set-name ,  
//DISP=(NEW,CATLG,DELETE) ,  
//UNIT=MAXSIZE ,  
//SPACE=( . . . .
```

would place the NEWTRANS data set on a volume in pool MAXSIZE (MAXSIZE must be a DASD pool). DASDPOOL FLST/RLST rules may then be used to change the pool or any other SET parameter.

DASDPOOL assignments override any JCL-specified UNIT and VOL=SER parameters.

DASDPOOL does not work with allocations made through SVC 32.

VSAM Component Separation

The data and index components of IDCAMS-defined key-sequenced VSAM data sets can be assigned to separate volumes by DASDPOOL through the use of the SEP rule list parameter. Both CANDIDATE=YES and SEP=YES or ASIS must be specified on the rule. The SEP parameter does not apply to JCL-defined VSAM data sets.

SEP=YES causes the data and index components to be assigned to separate volumes in the pool. If the number of volumes defined in the pool is insufficient, alternate pools are searched to find enough volumes to satisfy the request. If no alternate pools are defined, or none of the alternate pools has enough volumes, the allocation fails on the index component.

SEP=ASIS causes the data and index components to be assigned to separate volumes if the original definition of the VSAM cluster specifies different volumes for the two components. The usage of alternate pools is the same as for SEP=YES.

The exact effect of VSAM component separation depends on the way the VSAM data set is defined. Table 3-2 summarizes how the VSAM components are separated under various conditions.

Table 3-2 VSAM Component Separation Conditions

Option	Single Volume Definition	Multi-Volume Definition at the Cluster Level	Multi-Volume Definition at the Component Level
SEP=YES	Data/index separated within pool(s)—if pool has insufficient volumes, alternate pool used—if no alternate pool, index allocation fails	Multiple volumes assigned for both components—number same as in original cluster definition—(SEP has the effect of doubling volume requirements)—if pool has insufficient volumes, alternate pool used—if no alternate pool, index allocation fails	Number of volumes assigned matches that defined in the component definitions—if pool has insufficient volumes, alternate pool used—if no alternate pool, index allocation fails
SEP=NO	No separation processing performed—data/index are assigned to the same volume	No separation processing performed—data/index are assigned to the same volume	No separation processing performed—data/index are assigned to the same volume
SEP=ASIS	If a single volume is defined at cluster level, or the same single volume is defined at the component levels, no separation processing performed—if different single volumes are specified at the component level, processing is same as the Multi-Volume at Component Level column	No separation processing performed—data/index are assigned to the same volume	If the volume lists in both components are identical, processing is as if SEP=NO were specified—if the volume lists differ, processing is as if SEP=YES were specified

The use of alternate pools for VSAM component separation is highly recommended, particularly for multi-volume allocations.

VSAM component separation is not supported for JCL-defined VSAM data sets.

CANDIDATE parameter

The CANDIDATE parameter applies only to non-JCL defined VSAM data sets, for allocation of secondary volumes. If CANDIDATE=YES, DASDPOOL will allocate as many volumes from the selected pool as are in the IDCAMS DEFINE VOL=() list; the first volume selected from the pool will be allocated as the primary, while the other volumes from the pool will be candidate volumes. If CANDIDATE=NO is specified, the IDCAMS DEFINE VOL specification will be ignored, and there will be no secondary volumes available for normal OS/390 secondary volume allocation. (However, the MAINVIEW SRM function SPACVOLA can be used to allocate secondary volumes from the desired pool.) Note that with specification of the CANDIDATE parameter, the volume or volumes are assigned from the pool, not from those volumes specified in the IDCAMS DEFINE VOL list. Note also that CANDIDATE overrides NVOL. The CANDIDATE parameter does not apply to JCL-defined VSAM data sets; the number of volumes can be controlled using the NVOL and NVOLINDX parameters.

SUPVOL parameter

The SUPVOL parameter allows requests for specific volumes to be accepted or rejected. This enables the data center to direct data sets to specific volumes in certain situations, while disallowing it in general. When SUPVOL=NO is specified, a request for a specific volume (from JCL, IDCAMS DEFINE, and so on) is not suppressed; the specific volume requested will be assigned, and the data set will not be pooled. The rule may omit the pool name (SET SUPVOL=NO); however, if a pool name is specified (SET POOL=name SUPVOL=NO), the pool is ignored. When invoked, SUPVOL=NO processing will generate message number SMS3320I.

USEVOL parameter

The USEVOL parameter allows non-specific allocation requests to be directed to volumes based on use attribute (storage or private) as well as pool definition. This allows significant flexibility in volume assignment.

REJECT parameter

The REJECT parameter allows strict control of volume assignments by directing that the recommended device assignment (from OS/390) be rejected. If rejected, the data set allocation fails, and the job fails.

NVOLxxxx parameters

The NVOLxxxx parameters allow the number of volumes for a multi-volume allocation to be altered. NVOLMAX restricts the total volumes to the number of volumes in the assigned pool. NVOL and NVOLINDX restrict total volumes to the specified number for non-VSAM and VSAM data sets; both parameters are applied even if the data set is not assigned to a pool.

VFORCE parameter

The VFORCE parameter allows enforcement of standard VSAM component suffixes (DATA and INDEX). If VFORCE=YES, a DATA or INDEX suffix is added to non-JCL defined VSAM component names that do not already have that suffix. This parameter is applied regardless of pool assignment.

UNIT parameter

The UNIT parameter allows a unit name to be specified for a data set allocation. This parameter is applied regardless of pool assignment.

Warning! If the pool name (UNIT) is defined as a name that is coded in JCL, the function is active regardless of whether the mode is specified as SIMulate. All data sets that are included or not explicitly excluded from this rule for pooling are automatically allocated to this JCL-coded pool. Therefore, defining the UNIT in this manner should be used with extreme care.

VOLSEL parameter

The VOLSEL parameter specifies the type of volume selection to be performed (BESTFIT, CRITDSN, DPO, HISTDPO, MAXSPACE, or PERCENT), allowing a significant degree of control over the selection of the volume within the pool. VOLSEL is not required, and there is no default value; if not specified, OS/390 selects a volume from those in the DASDPOOL-assigned pool. The VOLSEL parameter takes effect only if at least one pool has been assigned.

USELIMIT parameter

The USELIMIT parameter specified for pools also affects pool and volume selection for non-DFSMS data sets. The USELIMIT parameter attempts to ensure that sufficient space is left on a volume to allow existing data sets to be extended with secondary extents. USELIMIT specifies a percentage utilization that should not be exceeded for a volume; if at all possible, additional primary allocations are not made on a volume if they would cause the total allocated space to exceed the USELIMIT. USELIMIT is specified in the pool definition in SMPOOLxx and applied during DASDPOOL VOLSEL processing and SPACVOLA processing.

Note: If VOLSEL=CRITDSN is used, USELIMIT is ignored.

The USELIMIT parameter on a pool will not prevent a pool assignment, even if a volume within the USELIMIT percentage cannot be found. In this case, the last volume found which would satisfy the primary allocation will be selected.

Although multiple pools may be specified on a rule (for example, SET POOL=(TEST,WORK,TEMP)), USELIMIT is applied only to a single pool at a time; that is, the first pool is examined for space, including USELIMIT considerations; if no allocation can be made to the first pool, the second pool is examined, and so on.

Messages

The DASDPOOL function issues one or more of the following messages:

SVM3310I *job,step,dd,dsn* UNITNAME CHANGED FROM *old unit* TO *new unit*

SVM3311I *job,step,dd,dsn* REJECT = *yes/no* SET

SVM3312I *job,step,dd,dsn* USEVOL = *volume* SET, PRIVATE VOLUMES WILL BE ELIGIBLE

SVM3319I *job,step,dd,dsn* NONSPECIFIC REQUEST ASSIGNED

SVM3320I *job,step,dd,dsn* SUPVOL=NO CODED - NO FURTHER DASDPOOL FUNCTIONS WILL BE PROCESSED

SVM3330I *job,step,dd,dsn* MAXIMUM NUMBER OF UNITS ALLOWED WILL BE LIMITED BY THE POOL SIZE

SVM3331I *job,step,dd,dsn* NUMBER OF UNITS FOR THE COMPONENT SET TO nn

SVM3332I *job,step,dd,dsn* NUMBER OF UNITS SET TO nn

SVM3333I *job,step,dd,dsn* CANDIDATE = YES SET, POOLED VOLUMES ARE ELIGIBLE

SVM3334I *job,step,dd,dsn* CANDIDATE = NO SET, POOLED VOLUMES ARE INELIGIBLE

SVM3343W *job,step,dd,dsn* VOLUME SELECTION CRITERIA *xxxxxxxx* IGNORED, NO POOLS ASSIGNED

SVM3335I *job,step,dd,dsn* VOLUME SELECTION CRITERIA SET TO *parameter*

SVM3336I *job,step,dd,dsn* VSAM SEPARATION SET TO *parameter*

SVM3337I *job,step,dd,dsn* ENFORCEMENT OF VSAM LOW-LEVEL QUALIFIER SET TO *.data or .index*

SVM3338I *job,step,dd,dsn* ELIGIBLE POOL NAMES SET TO: *pools*

SVM3339I *job,step,dd,dsn* DPO WINDOW SET TO *nnnn* SECONDS

SVM3353I *job,step,dd,dsn* CRITLIST = *crit_list*, CRITBIAS = *crit_bias*, CRITFAIL = *crit_fail*, CRITEMC = *crit_emc*

SVM3354I *job,step,dd,dsn* CRITLIST FAILED ALL ELIGIBLE VOLUMES. ALLOCATION REJECTED

SVM3361W *job,step,dd,dsnn* ERROR ON RAID API CALL. R15 = *rc*

SVM3362I *job,step,dd,dsn* RAIDDEVICE OF *raid_type* REQUIRED

SVM3390I *job,step,dd,dsn* DPO RESPONSE TIME OBJECTIVE SET TO *nnnn* MILLISECONDS

DASDPOOL issues message SVM3460I when a data set is assigned to a volume in a pool (through the VOLSEL parameter):

SVM3460I *job,step,dd,dsn vol/unit* VOL (*volume id*) SELECTED

Tip: To allow specific data sets to be allocated on a specific volume, use SUPVOL=NO in the rule list, or select those data sets in the filter list for inactive mode.

For example, to allow the storage administrator to place any data set on a specific volume, without regard to pooling, the following should be placed at the top of the rule list:

SET	SUPVOL=NO
INC	USER=STGADM/

SUPVOL=NO allows a JCL-specified volume to override the pool specification in the rule list for DASDPOOL. (Note that a pool assignment is not necessary.)

An alternate method is to place the following at the top of the filter list for DASDPOOL:

```
SET  MODE=INACT
INC  USER=STGADM/
```

Data sets for user STGADMxx will not invoke DASDPOOL, because the function's mode is set to INACTIVE for any user ID beginning with STGADM.

DASD Performance Optimization

The MAINVIEW SRM DPO facility is a specialized algorithm used during data set allocation to select a volume based on short-term historical performance data.

During I/O, the amount of time needed to complete the read or write operation is called the **response time**. The response time for a device can vary greatly due to the number of tasks competing for access to the volume and the channel path(s) to the volume. The DPO facility records device response times every 30 seconds and tables this performance data, maintaining 30 minutes worth of data in storage.

During allocation, the DPO facility scans the response time tables for the pooled devices over the indicated time interval; for example, for the last two minutes. The volume selected for allocation is the volume with the best response time. Alternatively, DPO can be used to select volumes based on response time objectives; for example, to select a volume that can provide a response time in the range of 20 to 40 milliseconds for some data sets or jobs and a volume that can provide a response time in the range of 40 to 60 milliseconds for other less response-critical data sets or jobs.

Within the MAINVIEW SRM ISPF interface, performance data can be displayed for each pool.

Starting DPO

The DPO facility is started and stopped automatically as part of the SVOS started task.

Resources

DPO maintains a table of performance data above the 16MB line in extended private storage. The table contains one entry for each DASD UCB for 30 minutes (sampled twice a minute); each entry is 16 bytes long. If the SVALLDPO task fails on start-up due to insufficient storage, the REGION JCL parameter should be increased for the loader job.

Usage

VOLSEL=DPO must be specified in the DASDPOOL rule for all data sets that are to use the DASD Performance Optimization allocation method.

The DPOWIND rule list parameter is used to specify the time interval over which the performance analysis is performed. For example, DPOWIND=120 indicates that the device with the best response time in the last 120 seconds should be selected for allocation. Since 30 minutes worth of data is retained, 1800 seconds is the logical upper limit for the DPOWIND operand.

Allocation based on response time objectives, rather than just the best response time, is specified in the DPORMIN and DPORMAX rule list parameters. These parameters set the minimum and maximum response times that are desired for the device to receive the allocation. For example, DPORMIN=20 and DPORMAX=40 will attempt to select a device that can provide a response time in the range of 20 to 40 milliseconds. If both DPORMIN and DPORMAX are zero, the best performing device is selected.

When multiple data sets are allocated in the same jobstep and some or all are using DPO allocation, the DPORSEP parameter is used to prevent all data sets from allocating to the same device, thus possibly degrading its performance significantly. The separation factor in DPORSEP is added to a previously used device's response time and the resulting figure is used in the selection process. This will artificially increase the device's response time, and cause subsequent allocations in the same job step to be allocated to another device with the desired response time characteristics.

In pools with fairly even performance statistics, the value of DPORSEP should generally be the average difference in response times on the different volumes. In pools with large differences in performance, such as pools that contain both cached and non-cached devices, allocations will concentrate on the fastest volumes, regardless of the separation factor.

In order to get a valid test of the DPO options, you should wait 30 minutes after you start DPO before you try to pool using DPO. It will take 30 minutes to collect the a full set of performance data. Any DPO-based allocations made before a full set of performance data has been collected may not reflect your actual selection algorithm.

Performance Data Display

The MAINVIEW SRM interface displays performance information for the devices in a pool. From the MPOOL view, type the **P** action line command to display performance data. The MVOLPER view (Figure 3-1) displays performance data (over the last 5 minute interval, maintained by DPO) for a selected pool, SMS pool, EMC physical device, or IBM physical device.

Figure 3-1 MVOLPER View

```

02JAN2001 13:29:57 ----- INFORMATION DISPLAY -----
COMMAND ==>
CURR WIN ==> 1          ALT WIN ==>
>W1 =MVOLPER=====SJSG=====02JAN2001==13:29:57==MVS RM=====16
C Volume UCB   Resp Queue  Pend  Disc  Conn  Busy  Serv  EXCPs # of SSI
- ----- Addr Time  Time  Time  Time  Time  Time  Time ----- UCBs Targ
SMS011 85A3   4.43  0.00  0.17  0.00  4.26  0.03  4.26    3    1 SJSG
SMS301 853E   4.43  0.00  0.21  0.00  4.22  0.03  4.22    3    1 SJSG
SMS002 83BD   4.42  0.00  0.12  0.08  4.22  0.04  4.30    3    1 SJSG
SMS303 8540   4.39  0.00  0.21  0.00  4.18  0.03  4.18    3    1 SJSG
SMS001 8355   4.38  0.00  0.12  0.08  4.18  0.03  4.26    3    1 SJSG
SMS004 8357   0.62  0.00  0.25  0.12  0.25  0.00  0.37    1    1 SJSG
SMS21B 8524   0.50  0.00  0.25  0.00  0.25  0.00  0.25    1    1 SJSG
SMS306 8543   0.50  0.00  0.25  0.00  0.25  0.00  0.25    1    1 SJSG
SMS305 8542   0.50  0.00  0.25  0.00  0.25  0.00  0.25    1    1 SJSG
SMS304 8541   0.50  0.00  0.25  0.00  0.25  0.00  0.25    1    1 SJSG
SMS25A 8358   0.50  0.00  0.25  0.00  0.25  0.00  0.25    1    1 SJSG
SMS003 8356   0.50  0.00  0.25  0.00  0.25  0.00  0.25    1    1 SJSG
SMS302 853F   0.49  0.00  0.12  0.12  0.25  0.00  0.37    1    1 SJSG
MSG00 8576   0.37  0.00  0.12  0.00  0.25  0.00  0.25    1    1 SJSG
SMSA00 85A2   0.00  0.00  0.00  0.00  0.00  0.00  0.00    0    1 SJSG
SMSA01 8579   0.00  0.00  0.00  0.00  0.00  0.00  0.00    0    1 SJSG

```

Historical DASD Performance Optimization

The MAINVIEW SRM HISTDPO facility allows you to pool data sets based on performance data gathered over a period of up to two weeks. Historical performance data is obtained using historical performance data collector in the BMC Software StorageGUARD product, which must be licensed and operational at your site.

Data is stored in a data space to decrease the access time required to obtain the performance data during allocation. The data is averaged and compared against the options you choose to determine the optimum volume.

Usage

The system parameter, HISTDAYS, specifies the number of days of data to retrieve from the StorageGUARD historical performance data collector and temporarily store in a data space. Values are 0-14 (7 or 14 recommended), which indicate the number of days of data to be retrieved. The default value is 0, which indicates that the data space is created, but no historical performance data is gathered and no HISTDPO pooling can be performed. If this value is changed, the historical performance data collector (SVSGP) must be stopped and restarted. For a valid HISTDPO test, do not specify a number of HISTDAYS higher than the total number of days collected in the historical performance data collector.

VOLSEL=HISTDPO must be specified in the DASDPOOL rule for all data sets that are to use the HISTDPO allocation method.

Starting time, ending time, and days of collection are specified in HDPOSTIM, HDPOETIM, and HDPODAYS respectively.

Allocation based on response time objectives, rather than just the best response time, is specified in the HDPORMIN and HDPORMAX rule list parameters. These parameters set the minimum and maximum response times that are desired for the device to receive the allocation. For example, HDPORMIN=20 and HDPORMAX=40 will attempt to select a device that can provide a response time in the range of 20 to 40 milliseconds. If both HDPORMIN and HDPORMAX are zero, the best performing device is selected.

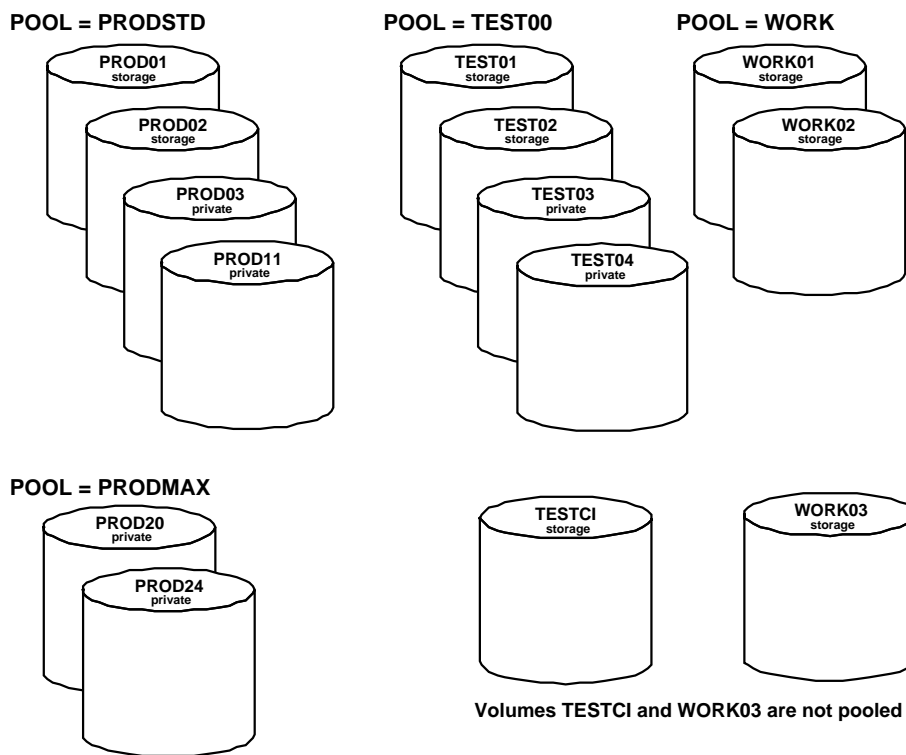
When multiple data sets are allocated in the same jobstep and some or all are using HISTDPO allocation, the HDPORSEP parameter is used to prevent all data sets from allocating to the same device, thus possibly degrading its performance significantly. The separation factor in HDPORSEP is added to a previously used device's response time and the resulting figure is used in the selection process. This will artificially increase the device's response time, and cause subsequent allocations in the same job step to be allocated to another device with the desired response time characteristics.

In pools with fairly even performance statistics, the value of HDPORSEP should generally be the average difference in response times on the different volumes. In pools with large differences in performance, such as pools that contain both cached and non-cached devices, allocations will concentrate on the fastest volumes, regardless of the separation factor.

Example

This example shows how to define pools for test, production, and work volumes, as shown in the following figure.

Figure 3-2 Possible System Storage Configuration

**SMMSYS10 member**

```
SET  FUNC=11 POOL=11 CAL=21
      VAR=11
      CMDPREF=# SMFID=250
      MSGID=YES TRKLEN=47476 TRKCYL=15
      PASSWORD=DB6145678901E
```

The MAINVIEW SRM system is defined.

SMFUNC11 member

```
SET  NAME=DASDPOOL ACTIVE=YES
      MSG=I SMF=I
      FLST=D4 RLST=D4
      DESC='ALLOCATE DATA SETS TO POOLS'
```

The function is defined and activated.

SMPOOL11 member

```
SET  POOLNAME=PRODSTD USELIMIT=90
      VOL=PROD0/
      VOL=PROD11
```

Pool PRODSTD includes all volumes whose names begin with PROD0 and volume PROD11. A 90% capacity limit is put on all volumes in this pool.

```
SET  POOLNAME=PRODMAX USELIMIT=80
      VOL=PROD2/
```

Pool PRODMAX includes two high-capacity 3390 volumes designated for very large data sets.

```
SET  POOLNAME=TEST00
      EXC VOL=TESTC/
      INC VOL=TEST/
```

Pool TEST00 includes all volumes whose names begin with TEST, except those beginning with TESTC.

```
SET  POOLNAME=WORK
      VOL=(WORK01,WORK02)
```

Pool WORK includes the two named volumes.

SMVARS11 member

```
SET  VARIABLE=SYSPGMRS
      VALUE=CRMORG/
      VALUE=SSMITH/
      VALUE=JAPPLEY/
      VALUE=DLITCHR/
      VALUE=AJOHNSN/
```

Define a variable that contains the user IDs of all system programmers.

SMFLSTD4 member

```
SET  MODE=INACT
      INC JOB=&SYSPGMRS DD=SYSDD/
```

By setting the mode to inactive, do not pool data sets with a DD name that starts with SYSDD that is in a job submitted by a system programmer. (This allows an override capability for system programmers; since pooling is not activated, the data set will be placed on the volume specified in the JCL.)

```
SET  MODE=ACT
      INC XMODE=JOB
      INC XMODE=TSO
```

For all other data sets allocated by batch jobs or TSO sessions, pooling is active.

SMRLSTD4 member

```
SET  SUPVOL=NO
      INC DSORG=VS VOL=DV5AR/
```

Do not pool any VSAM data sets that request allocation on volumes whose names begin with DV5AR.

```
SET  POOL=PRODMAX USEVOL=ALL
      VOLSEL=PERCENT
      EXC DSN=TEST/
      INC MAXSIZE>81MB
```

X Pool PRODMAX is selected for all data sets with a potential maximum size greater than 81MB, except for data sets whose name begins with TEST. All volumes in the pool (storage and private) are eligible for allocation.

BMC Software, Inc., Confidential and Proprietary Information

DSNCHECK—Enforce Data Set Naming Standards

The DSNCHECK function checks data set names for conformance to standards.

Data set name standards can significantly ease the transition to automated storage management. Standard naming conventions can embed important management information within the data set name, allowing more intelligent handling of data sets in storage management routines. The implementation of DFSMS is also easier if a standard data set naming process is in place.

DSNCHECK allows the data center to set and enforce standards for data set names. Name standards can be applied to the first eight data set name qualifiers and to the low-level qualifier. Using filter lists, rule lists, and name masking, flexible name standards can be specified and applied to resource groups selected on the basis of data set name, user name, job name, pool, volume, and so on.

Running DSNCHECK in simulation mode allows a gradual phase-in of new name standards. With normal messaging turned off (MSG=N) and SMF logging turned on (SMF=I), the administrator can preview the impact of new naming standards without any visible changes to the users. With normal messaging turned on (MSG=I), but still in simulation mode, the users are warned of impending standards implementation by the violation messages in their job logs (but without job cancellations). When an appropriate acclimatization period has passed, the function is put in active mode, and the naming standards are enforced.

System Parameters

The system parameters that affect DSNCHECK are described in Table 3-3.

Table 3-3 **System Parameters—DSNCHECK (Part 1 of 2)**

Parameter	Purpose
DASDGENR	specifies the only generic DASD names that are to be processed for specific functions if no space requirements are specified Standard pooling functionality allows the JCL UNIT parameter to specify a pool name. If this is not desired, DASDGENR can be used to intercept data set allocations without space information. Such allocations will only be considered for subsequent processing if the unit is found in an internal table (3380, 3350, and so on) or if the unit is found in the DASDGENR list. DASDGENR affects functions DASDPOOL, DSNCHECK, SETEXPDT, and FORCECAT.

Table 3-3 System Parameters—DSNCHECK (Part 2 of 2)

Parameter	Purpose
REJECT	controls termination of processing at the first or last data set rejected by REJECT=YES in DASDPOOL or DSNCHECK If termination is to take place on the first rejected data set, code REJECT=FIRST. If all data sets are to be processed before control is returned to OS/390 allocation, code REJECT=LAST.
TAPEGENR	specifies tape device generic names that some EasyPOOL functions will intercept If you want to intercept all tape requests, specify ALLTAPE as the first generic name. (However, you cannot use ALLTAPE when JCLEXT=NO.) TAPEGENR affects functions DSNCHECK, FORCECAT, SETEXPDT, and TAPEPOOL.

Rule List Parameter Quick Reference

Parameter	Description
DSNn=xxxxxxx	Name of data set qualifier 1–8
EVENTID=xxxxx	Identifies a user event defined in an SMEVNTxx member
LLQ=xxxxxxx	Name of low-level data set qualifier
MAXQLF=nn	Maximum number of data set name qualifiers
MINQLF=nn	Minimum number of data set name qualifiers
VCOMPLLQ=xxxxxxx	Name of VSAM component low-level data set qualifier

Rule List Parameter Explanations**DSNn=**

Purpose: Contains a data set name qualifier; *n* is a number from 1 to 8, indicating which data set name qualifier is to be checked.

If the data set name qualifier matches this parameter, the qualifier is accepted and the next qualifier or other parameter is checked; if the name qualifier does not match, the data set name is rejected with a naming violation (message SVM0344E) and no further name qualifiers are checked.

Note that while name masking cannot normally be used in a rule list parameter, in this particular case, name masking is accepted in the rule SET statement.

Syntax: DSN*n*=xxxxxxx

where *n* is a number in the range 1–8 and xxxxxx is any character string valid as a data set name qualifier.

Default: None

EVENTID=

Purpose: Specifies the identifier assigned to a user event in SMEVNTxx. This parameter will cause an event to be generated from this function.

Note: If EVENTID= is used on an FLST SET statement with MODE=INACT, the event will still be issued.

Syntax: EVENTID=xxxxxx

where xxxxx is the 5-character string specified on the EVNTID parameter in SMEVNTxx.

Default: None

LLQ=

Purpose: Contains a low-level data set name qualifier. If the last qualifier in the data set name matches this parameter, the qualifier is accepted; if the name qualifier does not match, the data set name is rejected with a naming violation (message SVM0344E).

Note that while MAINVIEW SRM name masking cannot normally be used in a rule list parameter, in this particular case, name masking is accepted in the rule SET statement.

Syntax: LLQ=xxxxxxxx

where xxxxxxxx is any character string valid as a data set name qualifier.

Default: None

MAXQLF=

Purpose: Specifies the maximum number of name qualifiers a data set is allowed. If the data set name contains more qualifiers than this number, the data set name is rejected with message SVM0342E.

Syntax: MAXQLF=nn

where nn is a number in the range 1–99.

Default: None

MINQLF=

Purpose: Specifies the minimum number of name qualifiers a data set is required to contain. If the data set name has fewer qualifiers than this number, the data set name is rejected with message SVM0341E.

Syntax: MINQLF=nn

where *nn* is a number in the range 1–99.

Default: None

VCOMPLLQ=

Purpose: Contains a low-level data set name qualifier for a VSAM data component. If the last qualifier in the VSAM component matches this parameter, the qualifier is accepted; if the name qualifier does not match, the data set name is rejected with a naming violation (message SVM0344E).

Note that while MAINVIEW SRM name masking cannot normally be used in a rule list parameter, in this particular case, name masking is accepted in the rule SET statement.

Syntax: VCOMPLLQ=xxxxxxx

where xxxxxxx is a valid data set name qualifier 1–8 characters long. It will contain blanks for the cluster.

Default: None

Parameters Not Supported

The following filter list parameters are not supported for DSNCHECK:

- GDGVER
- VALUE

Usage Notes

DSNCHECK provides services for VSAM and non-VSAM data sets. It operates on non-VSAM data sets at initial allocation and at data set rename. For VSAM data sets, it operates on DEFINE and ALTER operations for both the data and index components of the data set.

Note that DSNCHECK does not operate on:

- Data sets with only one qualifier in the name
- Temporary data sets

DSNCHECK acts on the first eight qualifiers in the data set name and on the low-level qualifier, regardless of how many qualifiers precede it.

Contrary to other rule list parameters, the DSNCHECK rule list parameters DSNn and LLQ can contain name masks. At execution time, these two parameters, as well as MINQLF and MAXQLF, are compared to their respective data set name components. If there is a match, the data set name qualifier is accepted; if there is not a match, a naming violation is reported. *Once a qualifier has failed a name check, no further qualifiers in the data set name are checked for compliance.* The MINQLF and MAXQLF parameters are checked first, followed by the DSNn parameters in numerical order, and then the LLQ parameter.

When DSNCHECK is activated for a rename operation, all DSN filter list parameters (DSN1-8, HQL, LLQ) contain values from the new data set name. A rename condition is identified by a value of RENAME in the ENVIR filter list parameter.

Naming standards violations are identified by one of three MAINVIEW SRM messages. If the function is running in active mode, the job is caused to fail with a JCL error (a VSAM IDCAMS DEFINE will fail with a VSAM catalog return code of 68 - unable to allocate space on user volume). For example:

```
06.24.50 JOB04488 SVM0344E EMPCRM , ,DD2,EMPCRM.V3R3M2.PROPLIB QUALIFIER(2) VIOLATES NAMING
STANDARDS
06.24.50 JOB04488 IEF452I EMPCRM - JOB NOT RUN - JCL ERROR
```

A name violation on a data set rename will be accompanied by OS/390 message “IEC614I - RENAME FAILED ...” for non-DFSMS data sets and by OS/390 message IEC331I for DFSMS data sets. Error messages for IEHPROGM or IDCAMS will also be produced with the SYSOUT for those utilities.

The VFORCE parameter of the DASDPOOL function also provides data set name services; VFORCE=YES adds standard component suffixes (DATA, INDEX) to VSAM file components.

Note that the DSNCHECK function must be specified in the SMFUNCxx member to be available and must be set active to provide enforcement services.

Naming standards and DFSMS are discussed in “Developing Naming Conventions” in the IBM® manual *MVS/ESA Storage Management Library: Managing Data*.

Example

This example shows how to:

- Require all load module partitioned data sets to end with LLIB, and all source partitioned data sets to end with SLIB.
- All partitioned data sets must have at least three qualifiers in their names.
- Run in simulate mode to assess the impact of this change.

SMFUNC11 member

```
SET  NAME=DSNCHECK ACTIVE=YES  
      MSG=I SMF=N  
      FLST=17 RLST=18  
      DESC='DATA SET NAME CHECKING'
```

The function is defined and activated. No SMF records will be produced.

SMFLST17 member

```
SET  MODE=SIM  
      EXC HLQ=SYS*  
      EXC HLQ=IS*  
      INC DSORG=PO
```

Data set name checking for resources selected by this filter list will be in simulate mode only. All partitioned data sets are selected, except those beginning with SYS and IS.

SMRLST18 member

```
SET  LLQ=*LLIB MINQLF=3  
      INC RECFM=U  
  
SET  LLQ=*SLIB MINQLF=3  
      INC LRECL=80
```

For all load libraries (RECFM=U), the low-level name qualifier must end with LLIB. Also, the data set name must contain at least three name qualifiers.

For all source libraries (record length of 80), the low-level name qualifier must end in SLIB, and the name must have at least three qualifiers.

Example

This example shows how to require the second qualifier of all data sets except system-level data sets (contained in pools SYSM1 and SYSRES) to be one of the following division codes: ATW, MKE, POS, HDA, CTY, FRP, with the following conditions:

- Data sets allocated by production jobs must have at least four name qualifiers, must have a valid DP system ID in the high-level qualifier, and must have a low-level qualifier beginning with P for non-VSAM data sets and PV for VSAM data sets.
- All other data sets must have at least three name qualifiers.
- PDSs must have LIB in the low-level qualifier.
- Data sets created by systems programmers are exempt from all naming standards, as are data sets residing on DB2® packs.

SMPOOLB2 member

SET	POOLNAME=SYSM1 TYPE=DASD INC VOL=SYS/ INC VOL=(BTD001,NRID0A,NRID33)	Define the SYSM1 pool to contain all volumes whose names start with SYS and three other volumes.
SET	POOLNAME=SYSRES INC VOL=*RES	Define the SYSRES pool to contain sysres volumes.
SET	POOLNAME=DB2POOL INC VOL=DB2/	Define the DB2POOL pool to contain all volumes whose names start with DB2.

SMFUNCB2 member

SET	NAME=DSNCHECK ACTIVE=YES MSG=I SMF=I FLST=B1 RLST=B1 DESC='ENFORCE NAME STANDARDS'	The function is defined and activated.
-----	---	--

SMVARSA1 member

SET	VAR=DIV VALUE=ATW VALUE=MKE VALUE=POS VALUE=HDA VALUE=CTY VALUE=FRP	Define a variable to hold all the division names.
SET	VAR=SYSPPOOL VALUE=SYSM1 VALUE=SYSRES	Define a variable to hold the two system pools.

SET	VAR=DPSYST VALUE=FIN* VALUE=HR* VALUE=MF* VALUE=DST*		Define a variable to hold the valid DP system name prefixes.
SMFLSTB1 member			
SET	MODE=ACT EXC USER=SYSP* EXC POOL=DB2POOL INC DSN=		All data sets are to be selected for processing, except those from users whose names start with SYSP, and those on a DB2 volume.
SMRLSTB1 member			
SET	MINQLF=2 INC POOL=&SYSPOOL		For all system-level data sets, the data set name must contain at least two qualifiers.
SET	DSN2=&DIV MINQLF=3 LLQ=*LIB* INC DSORG=PO		For partitioned data sets not in the system pools, the second qualifier must contain the division, the name must have at least three qualifiers, and the last qualifier must contain LIB.
SET	DSN1=&DPSYST DSN2=&DIV MINQLF=4 LLQ=PV* INC JOBCLASS=P DSORG=VS	X	For all VSAM data sets created by a production job, the first qualifier must contain a valid system ID, the second qualifier must contain the division, the name must have at least four qualifiers, and the last qualifier must start with PV.
SET	DSN1=&DPSYST DSN2=&DIV MINQLF=4 LLQ=P* INC JOBCLASS=P	X	For all non-VSAM data sets created by a production job, the first qualifier must contain a valid system ID, the second qualifier must contain the division, the name must have at least four qualifiers, and the last qualifier must start with P.
SET	DSN2=&DIV MINQLF=3 INC DSN=		For all other data sets, the second qualifier must contain the division, and the name must have at least three qualifiers.

FDRASIST—Pool Data Sets Recalled With FDR/ABR

The FDRASIST function assigns non-VSAM data sets being recalled by FDR/ABR to pools of DASD volumes. FDRASIST does not work for VSAM data sets.

Considerations in pooling during FDR recall include:

- Use of storage versus private volumes
- Specification of multiple pools, allowing overflow assignment if one or more pools are full
- Criteria for selection of the *best* volume in a pool

Along with the DASDPOOL and HSMRECAL functions, FDRASIST allows the data center to set and enforce standards for data sets' volume residency. Data sets being recalled by FDR can be directed to MAINVIEW SRM-defined pools on the basis of a variety of job and data set properties; for more flexibility in space utilization and management, multiple pools can be assigned, and multiple methods of volume selection can be used.

Rule List Parameter Quick Reference

Parameter	Description
CRITBIAS=	Defines the number of data sets that can reside on the volume before the current allocation
CRITEMC=YES/NO	Specifies whether the volume meeting CRITDSN criteria includes EMC physical volumes
CRITFAIL=YES/NO	Defines the allocation process if a volume meeting the criteria cannot be found
CRITLIST=xxxxxxx	Name of the table that contains allocation volumes
EVENTID=xxxxx	Identifies a user event defined in an SMEVNTxx member
POOL=xxxxxxx	Name of a pool 1–8 characters
USEVOL=STOR/PRIV/ALL	Type of volume
VOLSEL=BESTFIT/PERCENT/MAXSPACE	Method of volume selection from a pool

Rule List Parameter Explanations**CRITBIAS=**

Purpose: Defines the number of data sets that can reside on the volume before the current allocation.

Syntax: CRITBIAS=*nnn*

where *nnn* is a number from 1 to 999

Default: None

Note: This parameter works only in conjunction with
VOLSEL=CRITDSN

CRITEMC=

Purpose: Specifies whether the volume meeting CRITDSN criteria includes EMC physical volumes. If yes, data sets specified with CRITDSN should not reside on the same EMC Physical volume nor the same MVS Logical volume. If no, data sets specified with CRITDSN may reside on the same EMC physical volume. Default is no.

Syntax: CRITEMC=*YES/NO*

Default: CRITEMC=NO

CRITFAIL=

Purpose: Defines the allocation process if a volume meeting the criteria cannot be found. If YES, allocation fails; if NO, allocation is made to the best available volume. Like USELIM, if an acceptable volume cannot be found in the first pool, the best volume from the first pool is saved and the system tries the next pool(s). If no acceptable volume is found, CRITFAIL is processed on the saved volume.

Syntax: CRITFAIL=*YES/NO*

Default: CRITFAIL=NO

Note: This parameter works only in conjunction with
VOLSEL=CRITDSN

CRITLIST=

Purpose: The table that contains allocation volumes.

Syntax: CRITLIST=xxxxxxx

Default: None

Note: This parameter works only in conjunction with
VOLSEL=CRITDSN

Warning! CRITDSN is resource intensive and should be used only for a small list of critical data sets. It should not be used without considering the impact on the system.

EVENTID==

Purpose: Specifies the identifier assigned to a user event in SMEVNTxx. This parameter will cause an event to be generated from this function.

Note: If EVENTID= is used on an FLST SET statement with
MODE=INACT, the event will still be issued.

Syntax: EVENTID=xxxxx

where xxxxx is the 5-character string specified on the EVNTID parameter in SMEVNTxx.

Default: None

POOL=

Purpose: Specifies the name of a pool. Up to 15 values of this parameter can be specified within parentheses when used as a rule list parameter on a SET statement.

Syntax: POOL=xxxxxxx or POOL=(xxxxxxx,xxxxxxx,...)

where xxxxxxx is a valid pool name 1–8 characters long. If the poolname is not defined, the refresh or start-up will fail. Up to 15 pool names can be specified in parentheses.

Default: None

USEVOL=

Purpose: Specifies the type of volume (storage, private, or all) that satisfies non-specific allocation requests.

Syntax: USEVOL=xxxx

where xxxx is a value from the following list:

STOR	Storage mounted volume
PRIV	Privately mounted volume
ALL	Volume of any mount type

Default: USEVOL=ALL

VOLSEL=

Purpose: Specifies that volume selection within a pool should be based on current volume space availability.

Syntax: VOLSEL=BESTFIT/PERCENT/MAXSPACE

BESTFIT Specifies that the volume with the smallest contiguous extent that satisfies the primary allocation should be selected.

PERCENT Specifies that the volume with the largest amount of free space should be selected.

MAXSPACE Specifies that the volume with the largest single contiguous extent should be selected.

Default: VOLSEL=BESTFIT

Parameters Not Supported

The following filter list parameters are not supported for FDRASIST:

ACF2USER	BUFSP	CATALOG
CISIZE	CURDAY	CURSPACE
CURTIME	DD	DEVTYPE
DISP1	DISP2	DISP3
DSNTYPE	DYNALLOC	ENVIR
ERASE	FILESEQ	GDGVER
IMBED	JOBSDAY	JOBSTIME

LABELTYP	OLDDSN	OLDHLQ
OWNER	RELEASE	REPL
REUSE	UNIT	VALUE

Usage Notes

For FDRASIST to gain control, the volume requested on the NVOL parameter of the FDR SELECT statement must be rejected. This can be accomplished by specifying a volume in the FDR SELECT statement that does not exist in the pool assigned to the data set by the FDRASIST rule list. (The rejection also occurs if the NVOL-specified volume is full.)

Note: The NVOL-specified volume must exist and must be online.

FDRASIST will not work on VSAM data sets. For non-VSAM data sets, the FDR option ALCRSTIN in FDR releases 5.3 and higher, or the POOLDASD option in release 5.2, must be enabled. Refer to the FDR manual for additional information about this interface.

The module PLDRSTIN from the MAINVIEW SRM load library must be available at run time to the FDR job (for instance, in a steplib or the linklist).

The POOL parameter allows specification of one or more pools; pools are searched in the order specified. Note that any USELIMIT defined for a pool will be taken into account in selecting the pool and assigning a volume. However, USELIMIT will not prevent a pool assignment. If a pool that satisfies the uselimit requirement cannot be found, the last pool in the POOL parameter will be chosen.

The VOLSEL parameter specifies the type of volume selection to be performed, allowing a degree of control over the selection of a specific volume from a pool. If VOLSEL is not specified, the BESTFIT method will be used. Note that the MAXSPACE and PERCENT methods will always select a volume in a pool; however, if VOLSEL=BESTFIT cannot be satisfied, the pool/volume assignment will fail.

FDR does not interface with EasyPOOL for VSAM data sets. For VSAM data sets, DASDPOOL is used to assign data sets to pools of DASD volumes.

Example

This example shows you how to pool FDR/ABR recalls.

SMFUNC03 member

```
SET  NAME=FDRASIST ACTIVE=YES
      MSG=I SMF-I
      FLST=F1 RLST=F2
      DESC='POOL ON FDR RECALL'
```

SMPOOL10 member

```
SET  POOLNAME=FDR1
      VOL=WORK/

SET  POOLNAME=FDR2
      VOL=EMP/
```

SMFLSTF1 member

```
SET  MODE=ACT
      INC PGM=FDR*
```

FDRASIST is active for any program whose name begins with FDR.

SMRLSTF2 member

```
SET  POOL=FDR1
      INC PGM=FDR* HLQ=EMP

SET  POOL=FDR2
      VOLSEL=PERCENT
      INC PGM=FDR* HLQ=BBSD
```

Data sets with a high-level qualifier of EMP being processed by any program whose name begins with FDR will be recalled to pool FDR1, while data sets with a high-level qualifier of BBSD will be recalled to the volume with the largest single extent in pool FDR2.

FORCECAT—Force Catalog Disposition for New Data Sets

The FORCECAT function forces the normal termination disposition to CATLG for new data sets.

Effective storage management starts with properly tracking the location of data sets. Uncataloged data sets must be manually tracked or be lost, taking up valuable space on storage devices. Locating those uncataloged data sets may require a lengthy search of all volumes on the system. DFSMS manages only cataloged data sets. Data set cataloging is controlled by the JCL DD parameter DISP, but for non-DFSMS-managed data sets, OS/390 provides no way to require that the appropriate value of DISP is used.

For new data set allocations, FORCECAT forces the data set to be cataloged, regardless of the value of the DISP parameter in the JCL.

System Parameters

The system parameters that affect FORCECAT are described in Table 3-4.

Table 3-4 System Parameters—FORCECAT

Parameter	Purpose
DASDGENR	Specifies the only generic DASD names that are to be processed for specific functions if no space requirements are specified. Standard pooling functionality allows the JCL UNIT parameter to specify a pool name. If this is not desired, DASDGENR can be used to intercept data set allocations without space information. Such allocations will only be considered for subsequent processing if the unit is found in an internal table (3380, 3350, and so on) or if the unit is found in the DASDGENR list. DASDGENR affects functions DASDPOOL, DSNCHECK, SETEXPDT, and FORCECAT.
TAPEGENR	Specifies tape device generic names that some EasyPOOL functions will intercept. If you want to intercept all tape requests, specify ALLTAPE as the first generic name. (However, you cannot use ALLTAPE when JCLEXT=NO.) TAPEGENR affects functions DSNCHECK, FORCECAT, SETEXPDT, and TAPEPOOL.

Rule List Parameters

SET Statement (in member SMRLSTxx)

Parameter	Description
EVENTID=xxxxx	Identifies a user event defined in an SMEVNTxx member

EVENTID=

Purpose: Specifies the identifier assigned to a user event in SMEVNTxx. This parameter will cause an event to be generated from this function.

Note: If EVENTID= is used on an FLST SET statement with MODE=INACT, the event will still be issued.

Syntax: EVENTID=xxxxx

where xxxxx is the 5-character string specified on the EVNTID parameter in SMEVNTxx.

Default: None

Parameters Not Supported

The following filter list parameters are not supported for FORCECAT:

- GDGVER
- VALUE

Usage Notes

FORCECAT operates on any non-VSAM data set being created, DASD or tape, whose normal termination disposition is KEEP, for example, DISP=(xxx,KEEP). FORCECAT does not operate on DFSMS-managed data sets because DFSMS ensures that the data set is cataloged.

FORCECAT can be used in simulation mode to identify the new catalogs and aliases that must be established.

With the exception of the EVENTID rule list parameter, all resource selection must be done through the FLST.

Note that the FORCECAT function must be specified in the SMFUNCxx member to be available and must be set active to provide services.

FORCECAT generates message SVM3445I when it changes the normal termination disposition from KEEP to CATLG:

SVM3445I *job,step,dd,dsn* CATALOGING FORCED

Example

This example shows how to force cataloging for all data sets except those on volume NOCAT or tape pool EXPORT.

SMFUNCxx member

```
SET  NAME=FORCECAT ACTIVE=YES
      MSG=I SMF=I
      FLST=FC
      DESC='FORCE DATA SET
            CATALOGING'
```

The function is defined and activated. There is no RLST specification.

SMFLSTFC member

```
SET  MODE=ACT
      EXC VOL=NOCAT
      EXC POOL=EXPORT
      INC DSN=/'
```

This filter list specification is active. All data sets except those on NOCAT or pool EXPORT are selected.

Example

This example shows how to force data set cataloging for medium and large data sets created in batch and TSO jobs.

SMFUNCxx member

```
SET  NAME=FORCECAT ACTIVE=YES
      MSG=I SMF=I
      FLST=48
      DESC='FORCE DATA SET CATALOGING'
```

The function is defined and activated.

SMFLST48 member

```
SET  MODE=ACT
      EXC MAXSIZE<5MB
      INC XMODE=JOB
      INC XMODE=TSO
```

This filter list specification is active. It selects jobs that are run in batch or from TSO which have a potential maximum size greater than 5 megabytes.

MODDELET—Remove Specific Volume for MOD-DELETE Data Set Deletion

The MODDELET function removes a specific volume assignment on a data set deleted by DISP=(MOD,DELETE,DELETE).

A convenient and widely used method of deleting possibly-nonexistent data sets is specification of DISP=(MOD,DELETE,DELETE) in JCL. If the data set exists, it is found in the catalog and deleted. If the data set does not exist, it is created and then deleted. In either case, a data set which may or may not exist will be deleted, with no JCL errors. In some situations, however, a specific volume assignment is included in the MOD-DELETE job step. With MAINVIEW SRM pooling active, the data set to be deleted may not be on that volume, resulting in the failure of the data set deletion.

The MODDELET function removes the specific volume assignment from job steps using the MOD-DELETE data set deletion method, allowing the operation to proceed normally.

Rule List Parameters

SET Statement (in member SMRLSTxx)

Parameter	Description
EVENTID=xxxxx	Identifies a user event defined in an SMEVNTxx member

EVENTID=

Purpose: Specifies the identifier assigned to a user event in SMEVNTxx. This parameter will cause an event to be generated from this function.

Note: If EVENTID= is used on an FLST SET statement with MODE=INACT, the event will still be issued.

Syntax: EVENTID=xxxxx

where xxxxx is the 5-character string specified on the EVNTID parameter in SMEVNTxx.

Default: None

Parameters Not Supported

The following filter list parameters are not supported for MODDELET:

- GDGVER
- VALUE

Usage Notes

MODDELET affects only data sets in job steps where the executing program is IEFBR14 and DISP=(MOD,DELETE,DELETE) is specified. Any specific volume assignment (including multiple volumes) is removed from the allocation; in effect, VOL=SER=(volid,volid,. . .) is deleted from the JCL. The data set is then accessed from the system catalog.

MODDELET does not affect DFSMS-managed data sets or data sets that are allocated to VIO.

MODDELET generates message SVM3482I when it processes a MOD-DELETE action, including those without a specific volume assignment:

SVM3482I *job,step,dd,dsn* HANDLED BY MODDELET

Example

Remove specific volume assignments from MOD-DELETE data set deletions.

SMFUNCxx member

```
SET  NAME=MODDELET ACTIVE=YES
      MSG=I SMF=I
      FLST=M4
      DESC='PROCESS DISP=(MOD,DELETE)'
```

The function is defined and activated. There is no RLST specification.

SMFLSTM4 member

```
SET  MODE=ACT
      INC JOBCLASS=P
```

This filter list specification is active. All data sets in a production job class are selected for processing.

OPENEMPT—Initialize Empty Data Sets

The OPENEMPT function initializes empty data sets at job step termination.

Data sets that are allocated but never opened are a potential hazard to subsequent jobsteps and jobs. OS/390 does not initialize a data set at allocation; a new data set is initialized only when it is opened. If an uninitialized data set is later opened by another program, one of three occurrences will happen:

- The job will receive an I/O error because an end-of-file mark has not been written.
- The job will read a record left on the volume from a prior data set and possibly process it as normal data.
- The System-Determined Blocksize (SDB) may not be driven, resulting in an S001 abend.

OPENEMPT can intercept and prevent most uninitialized data set problems by initializing empty data sets at job step termination.

Rule List Parameters

SET Statement (in member SMRLSTxx)

Parameter	Description
EVENTID=xxxxx	Identifies a user event defined in an SMEVNTxx member

EVENTID=

Purpose: Specifies the identifier assigned to a user event in SMEVNTxx. This parameter will cause an event to be generated from this function.

Note: If EVENTID= is used on an FLST SET statement with MODE=INACT, the event will still be issued.

Syntax: EVENTID=xxxxx

where xxxxx is the 5-character string specified on the EVNTID parameter in SMEVNTxx.

Default: None

Parameters Not Supported

The following filter list parameters are not supported for OPENEMPT:

BUFSP	CATALOG	CISIZE
CURDAY	CURSPACE	CURTIME
ERASE	EXTENT	GDGVER
IMBED	JOBSDAY	JOBSTIME
LABELTYP	OWNER	REPL
REUSE	VALUE	

Usage Notes

OPENEMPT provides services for new non-VSAM data sets other than partitioned data sets. OPENEMPT obtains control at the termination of the jobstep and checks all data sets for an uninitialized condition. If found, the uninitialized data sets are opened and closed to initialize them.

An additional benefit of OPENEMPT is that the initialized data set can be processed by the SPACRLSE function to release allocated unused space. If the data set remained uninitialized, it would retain its full allocated space.

OPENEMPT ignores the data set in the following circumstances:

- The step is terminating normally, and the normal termination disposition is delete: DISP=(NEW,DELETE)
- The step is terminating abnormally, and the abnormal termination disposition is delete; DISP=(NEW,anything,DELETE)
- The data set has no volumes assigned
- The data set is a PDS
- The data set is subsequent to the first data set in a concatenation of data sets (only the first data set is processed due to OS/390 architectural restrictions)
- The data set is allocated to VIO
- No DCB information specified; DCB information uses a referback (DCB=*.STEP.DD)

Messages

The messages issued by OPENEMPT are:

```
SVM0500I  dsn INITIALIZED
SVM0501I  dsn NOT INITIALIZED. REASON = reason
SVM0500I  job,step,dd,dsn INITIALIZED
```

OPENEMPT found the indicated data set had not been initialized. The data set was opened and closed by OPENEMPT to set the end-of-file mark.

```
SVM0501I  job,step,dd,dsn NOT INITIALIZED, REASON=RECFM
```

The indicated data set could not be initialized because record format information was not available; for example:

```
SVM0501I  EMPCRM,STEP10,DD2,EMPCRM.OPENEMPT.TEST NOT
INITIALIZED,REASON=RECFM
```

```
//EMPCRM JOB (3500), 'MORRIS',MSGCLASS=R,CLASS=F
//STEP10 EXEC PGM=IEFBR14
//DD2 DD DSN=EMPCRM.OPENEMPT.TEST,
// DISP=(NEW,CATLG,DELETE),
// SPACE=(TRK,(1,1)),UNIT=SYSDA
```

```
SVM0501I  job,step,dd,dsn NOT INITIALIZED, REASON=DDNAME
```

The indicated data set could not be initialized because the ddname was a duplicate of another; for example:

```
SVM0500I  EMPHRFX,STEP10,DD2,EMPHRF.OPENEMPT.TEST INITIALIZED
SVM0501I  EMPHRFX,STEP10,DD2,EMPHRF.OPENEMPT.TEST2 NOT
INITIALIZED,REASON=DDNAME
```

```
//EMPHRFX JOB (3500), 'FITZGERALD',MSGCLASS=R,CLASS=F
//STEP10 EXEC PGM=IEFBR14
//DD2 DD DSN=EMPHRF.OPENEMPT.TEST,
// DISP=(NEW,CATLG,DELETE),
// SPACE=(TRK,(1,1)),UNIT=SYSDA,
// DCB=(LRECL=80,BLKSIZE=16000,RECFM=FB)
//DD2 DD DSN=EMPHRF.OPENEMPT.TEST2,
// DISP=(NEW,CATLG,DELETE),
// SPACE=(TRK,(1,1)),UNIT=SYSDA,
// DCB=(LRECL=80,BLKSIZE=16000,RECFM=FB)
```

```
SVM0501I  job,step,dd,dsn NOT INITIALIZED, REASON=OPEN
```

The indicated data set could not be initialized because it was already open. This normally indicates the program terminated without closing the data set.

Note that the OPENEMPT function must be specified in the SMFUNCxx member to be available and must be set active to provide services.

Example

This example shows how to use OPENEMPT on all data sets other than those in programs under development.

SMFUNC11 member

SET	NAME=OPENEMPT	ACTIVE=YES	The function is defined and activated. Note that there is no RLST parameter.
	MSG=I	SMF=I	
	FLST=03		
	DESC='OPEN EMPTY DATA SETS'		

SMVARS02 member

SET	VARIABLE=TESTPGMS	Set a variable to identify test jobs by program name. Any job that begins or ends with T, or that begins with DBT is considered to be a test job.
	VALUE=T*	
	VALUE=*T	
	VALUE=DBT*	
SET	VARIABLE=TESTJOBS	Set a variable to identify test jobs by job class.
	VALUE=T	
	VALUE=D	
	VALUE=E	

SMFLST03 member

SET	MODE=ACT	Select resources for OPENEMPT based on program name or job class.
INC	PGM=&TESTPGMS	
INC	JOBCLASS=&TESTJOBS	

OPTBLKSZ—Set Optimum Blocksize

The OPTBLKSZ function forces an optimum blocksize for new data set allocations on DASD and tape. With the use of optional parameters, the BLKSIZE may also be changed on OLD or INPUT data sets.

The blocksize of a data set has a direct effect on the storage efficiency of DASD devices. Generally, the larger the blocksize, the more records fit on a track. This not only achieves better utilization of disk space, it also reduces the number of I/Os required to read the data set. Although very small data sets may not benefit significantly from the denser record storage, the minimum allocation by OS/390 is a single track; therefore, it generally is more effective to allocate larger blocksizes. The only disadvantage is that larger blocksizes require larger buffers in memory. The following table shows the effect of different blocksizes on DASD capacity.

Record Density, 3380 DASD, 80-byte records			
Blocksize	Records/track	Records/cylinder	Percent space used ^a
4240	530	7950	90.0
6320	553	8295	93.7
23440	586	8790	98.9

^a Source: IBM 3380 Direct Access Storage: Reference Summary

OPTBLKSZ allows the data center to set and enforce standard data set blocksizes. Using rule lists, different blocksizes can be applied to different groups of data sets.

System Parameters

The system parameters that affect OPTBLKSZ are described in Table 3-5.

Table 3-5 System Parameters—OPTBLKSZ (Part 1 of 2)

Parameter	Purpose
BLKINPUT	Changes block size for input data sets.
BLKOLDSR	Changes blocksize for output data sets opened with disposition of old or shared.
TRKCYL	Specifies the number of tracks per cylinder for the default device type. The value specified for 3380/3390/9345 devices should be fifteen. (Note that this specification is the same as the SCDS base configuration DEFINE under ISMF for DFSMS.) TRKCYL and TRKLEN are used by the DASDPOOL function to convert allocations in tracks or cylinders to megabytes for volume selection based on available space; for example, VOLSEL=BESTFIT. The information specified on these two parameters should reflect the devices that are most prevalent in your environment.

Table 3-5 System Parameters—OPTBLKSZ (Part 2 of 2)

Parameter	Purpose
TRKLEN	<p>Specifies the number of bytes per track for the default device type. Valid values are: 3380 – 47,476 3390 – 56,664 9345 – 46,456</p> <p>Note that this specification is the same as the SCDS base configuration DEFINE under ISMF for DFSMS.)</p> <p>TRKCYL and TRKLEN are used by the DASDPOOL function to convert allocations in tracks or cylinders to megabytes for volume selection based on available space; for example, VOLSEL=BESTFIT. The information specified on these two parameters should reflect the devices that are most prevalent in your environment.</p>

Rule List Parameters

SET Statement (in member SMRLSTxx)

Parameter	Description
BLKSIZE=nnnnn	A blocksize in the range 0–32,760
EVENTID=xxxxx	Identifies a user event defined in an SMEVNTxx member
FORCE=YES/NO	Overrides program specified blocksize

Rule List Parameter Explanations**BLKSIZE=**

Purpose: Specifies a blocksize to be assigned to the new data set. A specification of zero directs the use of a system-assigned blocksize. This is a required parameter; no processing occurs if a BLKSIZE is not specified.

Syntax: BLKSIZE=nnnnn

where nnnnn is a number in the range 0–32,760.

Default: None

EVENTID=

Purpose: Specifies the identifier assigned to a user event in SMEVNTxx. This parameter will cause an event to be generated from this function.

Note: If EVENTID= is used on an FLST SET statement with MODE=INACT, the event will still be issued.

Syntax: EVENTID=xxxxx

where *xxxxx* is the 5-character string specified on the EVNTID parameter in SMEVNT*xx*.

Default: None

FORCE=

Purpose: Specifies whether any program specified blocksize should be overridden.

Syntax: FORCE=*YES/NO*

If YES, the value specified will be used to override a program specified blocksize; if NO, a program specified blocksize will NOT be overridden.

Default: FORCE=NO

Parameters Not Supported

The following filter list parameters are not supported for OPTBLKSZ:

- GDGVER
- VALUE

Usage Notes

OPTBLKSZ provides services for both new and old DASD and tape data sets with record formats of undefined, variable, fixed, or fixed blocked. *It does not operate on ISAM, direct, or VSAM data sets.* It will override blocksizes assigned in program DCBs if FORCE=YES is specified. The data set allocation *must* be specified with both record length and a blocked record format (FB, VB, FBA, and so on); without these two pieces of information, a blocksize is not assigned.

A BLKSIZE rule list parameter is required; if not specified in a rule, the data set selected for processing is ignored.

If the LRECL of the selected data set is zero, or larger than the rule-specified blocksize (SET BLKSIZE=), OPTBLKSZ takes no action. Otherwise, a new blocksize is assigned. For fixed and fixed blocked data sets, the rule-specified blocksize is adjusted to the next lower multiple of the LRECL; for variable and undefined data sets, the rule-specified blocksize is used as is.

Note that a blocksize of zero can be specified to request a system-assigned blocksize; in this case, OPTBLKSZ sets the BLKSIZE to 0. OS/390 System-Determined Blocksize (SDB) must be driven to reset the BLKSIZE to the optimum value. For SDB to reset the value, the data set must be opened in the same jobstep in which it is allocated (see “OPENEMPT—Initialize Empty Data Sets” on page 3-50).

Messages

The message issued by OPTBLKSZ when a blocksize is changed is:

SVM3347I *job,step,dd,dsn* BLKSIZE CHANGED *nnnnn/nnnnn*

The first *nnnnn* gives the original blocksize specified in the JCL; the second *nnnnn* gives the blocksize set by OPTBLKSZ.

The message issued by OPTBLKSZ when a specific blocksize is requested by program name and FORCE=YES is coded is:

SVM3351I *job,step,dd,dsn* BLKSIZE *nnnn* IN PGM=*name* REMOVED

Note that the OPTBLKSZ function must be specified in the SMFUNCxx member to be available and must be set active to provide service.

Example

This examples shows you how to:

- Specify blocksizes for all data sets except certain system-level PDSs.
- Run in simulate mode for data sets in jobs with BLKEXEMPT in the third job account field.

SMFUNC73 member

```
SET  NAME=OPTBLKSZ ACTIVE=YES
      MSG=I SMF=N
      FLST=7D RLST=7E
      DESC='ASSIGN OPTIMUM BLOCKSIZE'
```

The function is defined and activated. No SMF records will be produced.

SMFLST7D member

```
SET  MODE=SIM
      INC JOBACCT3=BLKEXEMPT
```

The function is to run in simulate mode for all data sets in jobs with BLKEXEMPT in the third job accounting field.

```
SET  MODE=ACT
      EXC HLQ=SYS* DSORG=PO RECFM=U
      INC DSN=/'
```

The function is to run in active mode for all selected resources. All data sets are selected, except system-level undefined partitioned data sets.

SMRLST7E member

```
SET  BLKSIZE=3861
      INC DSORG=PO LLQ=JCL/'
```

Specify a small blocksize for JCL PDSs, since they consist of many members with 50 records or less.

```
SET  BLKSIZE=23440
      INC MAXSIZE>50MB DEVTYPE=3380
      INC MAXSIZE>50MB DEVTYPE=3390
```

Specify half-track blocking for large data sets on all 3380 and 3390 devices.

```
SET  BLKSIZE=4629
      INC MAXSIZE>50MB DEVTYPE=DASD
```

For large data sets on all other devices, the blocksize is set for half- or third-track blocking.

```
SET  BLKSIZE=11476
      INC DEVTYPE=3380
```

For all other data sets on 3380 devices, the blocksize is set to quarter-track blocking.

```
SET  BLKSIZE=13682
      INC DEVTYPE=3390
```

For all other data sets on 3390 devices, the blocksize is set to quarter-track blocking.

```
SET  BLKSIZE=25600
      INC DEVTYPE=TAPE
```

The blocksize for tape data sets is set to 25KB.

```
SET  BLKSIZE=27999 FORCE=YES
      INC PROGRAM=PRO031/'
```

For all programs that start with PRO031, set the blocksize to 27999 and override any program blocksize.

SETEXPDT—Set Expiration Date for Data Set

The SETEXPDT function sets expiration dates for data sets being created.

The OS/390 expiration date for a data set prevents deletion of that data set by normal processing; that is, by disposition processing specified in a JCL DD statement, or by deletion through ISPF services. Proper use of the expiration date facility protects data sets from deletion while they are still useful; improper use of expiration dates can retain data sets long past their useful life, resulting in wasted DASD space and tape volumes.

SETEXPDT allows the data center to set standards for retention periods and to enforce those standards rigorously. Application of expiration dates can be based on data set names, on device types, on assigned pools, on size, or on any property allowed as a MAINVIEW SRM filter list parameter. In addition, the calculation of an expiration date can take into account the non-working days that occur in the retention period.

SETEXPDT also provides a convenient method of assigning symbolic expiration dates to selected data sets; for example, assigning CA-1 special-use dates (out-of-library, days since last used, and so on). New releases of CA-1 use ACCODE to determine special use tapes. The AC-CODE SET parameter supports this feature.

System Parameters

The system parameters that affect SETEXPDT are described in Table 3-6.

Table 3-6 **System Parameters—SETEXPDT**

Parameter	Purpose
DASDGENR	Specifies the only generic DASD names that are to be processed for specific functions if no space requirements are specified. Standard pooling functionality allows the JCL UNIT parameter to specify a pool name. If this is not desired, DASDGENR can be used to intercept data set allocations without space information. Such allocations will only be considered for subsequent processing if the unit is found in an internal table (3380, 3350, and so on) or if the unit is found in the DASDGENR list. DASDGENR affects functions DASDPOOL, DSNCHECK, SETEXPDT, and FORCECAT.
TAPEGENR	Specifies tape device generic names that some EasyPOOL functions will intercept. If you want to intercept all tape requests, specify ALLTAPE as the first generic name. (However, you cannot use ALLTAPE when JCLEXT=NO.) TAPEGENR affects functions DSNCHECK, FORCECAT, SETEXPDT, and TAPEPOOL.

Rule List Parameters

SET Statement (in member SMRLSTxx)

Parameter	Description
AC_CODE=xxxxxxx	Specifies the value in the IBM ACCODE field.
CAL=YES/NO	Adjusts calculated expiration date by the non-working days calendar
EVENTID=xxxxx	Identifies a user event defined in an SMEVNTxx member
EXPDT=yyddd or yyyddd	New expiration date in Julian format
REPLACE=YES/NO	Allows override to JCL-specified expiration date
RETPD=nnnn	Retention period in days

Rule List Parameter Explanations

AC_CODE=

Purpose: Specifies the value in the IBM ACCODE field. Normally, this field is used in conjunction with user-generated tape labels. CA's tape manager, CA-1, uses the field to indicate various special tapes, such as off-site, permanent hold, and so on.

Syntax: AC_CODE=xxxxxxx

Default: None

CAL=

Purpose: Specifies that the calculation of a new expiration date must consider non-working days in MAINVIEW SRM's non-working day calendar. If EXPDT is specified, CAL is ignored, since EXPDT directs a specific date. However, if RETPD is specified with CAL=YES, the new expiration date is calculated by adding the retention period and the number of non-working days in that period to the creation date.

Syntax: CAL=YES/NO

Default: CAL=YES

Note: ISPF function 3.4 does not correctly display expiration dates for JCL-defined VSAM data and index component expiration dates. LISTCAT does show the correct dates.

EVENTID=

Purpose: Specifies the identifier assigned to a user event in SMEVNTxx. This parameter will cause an event to be generated from this function.

Note: If EVENTID= is used on an FLST SET statement with MODE=INACT, the event will still be issued.

Syntax: EVENTID=xxxxxx

where xxxxx is the 5-character string specified on the EVNTID parameter in SMEVNTxx.

Default: None

EXPDT=

Purpose: Specifies a data set expiration date. Both CAL and RETPD are ignored if EXPDT is specified.

The specified Julian date is not checked for leap year validity; that is, if you specify a non-leap-year with 366 days (for example, 97366), it is accepted.

Syntax: EXPDT=nnnnn

where nnnnn is a Julian date in the format yyddd to represent a date from 1900–1999.

OR

EXPDT=nnnnnnnn

where nnnnnnn is a Julian date in the format yyyyddd to represent a date from any year.

Default: None

REPLACE=

Purpose: Specifies whether an expiration date already specified for the data set (using JCL, for example) is overridden. If REPLACE=NO and an expiration date already exists for a data set, then the SETEXPDT function does not replace it. If REPLACE=YES, the SETEXPDT function sets the expiration date, regardless of any existing expiration date.

Syntax: REPLACE=YES/NO

Default: REPLACE=YES

RETPD=

Purpose: Specifies the retention period in days for a new data set. If CAL=YES is also specified, the retention period **and** any non-working days within that period are added to the data set's creation date to calculate the new expiration date. If CAL=NO is specified (or defaulted), the retention period is added to the data set's creation date to calculate the new expiration date; non-working days are ignored.

Note that a specification of EXPDT in the same SET statement overrides any RETPD specification.

Syntax: RETPD=nnnn

where *nnnn* is a number in the range 0–9999.

Default: None

Parameters Not Supported

The following filter list parameters are not supported for SETEXPDT:

- GDGVER
- VALUE

Usage Notes

SETEXPDT provides services only for data sets being created; data sets that already exist are not changed by SETEXPDT. For VSAM data sets, SETEXPDT works for JCL-defined VSAM but not for IDCAMS-defined VSAM.

Specification of the expiration date parameter (EXPDT) directs the usage of a specific date and overrides any specification of retention period (RETPD) and non-working days (CAL). If EXPDT is omitted and RETPD specified, the new expiration date is calculated by adding the retention period (in days) to the data set's creation date. If CAL=YES is also specified, the non-working days in the retention period (specified in the SMCALSxx member) are also added to the creation date.

If neither EXPDT nor RETPD is specified, no expiration date is set by the function, even though a data set may satisfy all filter list and rule list requirements.

The REPLACE parameter allows SETEXPDT to be applied in a restrictive or a lenient manner. The REPLACE parameter specifies whether a previously-specified expiration date is overridden by SETEXPDT. REPLACE=YES forces the use of the expiration date established in SETEXPDT, regardless of any expiration date or retention period already set. REPLACE=NO, however, accepts the original expiration date specification (in JCL, for example); the expiration date calculated by SETEXPDT is applied only if there is no existing expiration date specification.

Note that the SETEXPDT function must be specified in the SMFUNCxx member to be available and must be set active to provide services.

Messages

SETEXPDT generates message SVM3453I when it provides an expiration date for a data set:

SVM3453I *job,step,dd,dsn* EXPDT (yyddd) ASSIGNED

SETEXPDT generates message SVM3348I when the ACCODE field for the current allocation has been set to *ac_code*.

SVM3348I *job,step,dd,dsn* ACCODE SET TO *ac_code*

Example

This example shows you how to set an expiration date for all data sets created by a batch job in the pool PROD, based on a retention period of 180 days. It overrides any expiration date already specified for the data set. You must take non-working days into account.

SMFUNC_{xx} member

SET	NAME=SETEXPDT	ACTIVE=YES	The function is defined and activated.
	MSG=I	SMF=I	
	FLST=64	RLST=73	
	DESC=' SET EXPDT FOR DASD & TAPE'		

SMFLST64 member

SET	MODE=ACT	This filter list specification is active. It selects all data sets in pool PROD.
	INC P00L=PROD	

SMRLST73 member

SET	RETPD=180	REPLACE=YES	CAL=YES	For all batch jobs (XMODE=JOB), the retention period is 180 days, plus the number of non-working days in that 180-day period.
	INC	XMODE=JOB		

SET	RETPD=10	REPLACE=YES	Allow a maximum retention period of 10 days for data sets created by test jobs.	
	INC	RETPD>10	JOBCLASS=T	

Example

This example shows you, for all data sets in pool PROD, how to:

- Set a specific expiration date for GL93 data sets whose names do not end with TEST, if created by a batch job and if an expiration date has not been specified.
- Set a 30-day retention period for those data sets whose names end with TEST. Require the use of this new expiration in place of any expiration date already specified for the data set.

SMFUNCxx member

```
SET  NAME=SETEXPDT ACTIVE=YES
      MSG=I SMF=I
      FLST=5A RLST=7A
      DESC='SET EXPDT FOR DASD & TAPE'
```

The function is defined and activated.

SMFLST5A member

```
SET  MODE=ACT
      INC POOL=PROD
```

This filter list specification is active. It selects all data sets in pool PROD.

SMRLST7A member

```
SET  EXPDT=94365 REPLACE=NO
      INC XMODE=JOB DSN=GL93/ LLQ-=TEST
```

For all data sets created by a batch job, with a name starting with GL93 and ending with something other than TEST, and which do not already have an expiration date, an expiration date of the last day in 1994 is specified.

```
SET  RETPD=30 REPLACE=YES
      INC DSN=GL93*.*.TEST
```

For all data sets starting with GL93 and ending with TEST, a 30-day retention period is forced.

SMSACSTE—Test ACS Routines

The SMSACSTE function provides services that enhance the process of testing ACS routines:

- ACS variables are displayed, allowing easier debugging of ACS routines.
- The data set can be removed from DFSMS-management, allowing testing of ACS routines without generating numerous new DFSMS-managed data sets.

Rule List Parameters

SET Statement (in member SMRLSTxx)

Parameter	Description
EVENTID=xxxxx	Identifies a user event defined in an SMEVNTxx member
SMSMANAGED=YES/NO	Specifies whether the data set is to remain DFSMS-managed

Rule List Parameter Explanations

EVENTID=

Purpose: Specifies the identifier assigned to a user event in SMEVNTxx. This parameter will cause an event to be generated from this function.

Note: If EVENTID= is used on an FLST SET statement with MODE=INACT, the event will still be issued.

Syntax: EVENTID=xxxxx

where xxxxx is the 5-character string specified on the EVNTID parameter in SMEVNTxx.

Default: None

SMSMANAGED=

Purpose: Specifies whether the data set is to remain DFSMS-managed. NO removes the data set from DFSMS control; YES leaves the data set under DFSMS management.

Syntax: SMSMANAGED=YES/NO

Default: SMSMANAGED=YES

Parameters Not Supported

The following filter list parameters are not supported for SMSACSTE:

- GDGVER
- VALUE

Usage Notes

SMSACSTE applies only to data sets that are DFSMS-managed. The function displays most ACS variables and also allows the data set to be removed from DFSMS management. SMSACSTE is activated for data sets with a status disposition of NEW or MOD.

SMSACSTE recognizes a DFSMS-managed data set by the assignment of a storage class to the data set. If the data set is not DFSMS-managed, SMSACSTE does not process the data set. No messages are produced, the SMSMANAGED parameter is ignored and the trace facility shows no activity. Regardless of any construct assignments by SMSACSDC, SMSACSMC, or SMSACSSG, the SMSACSTE function does not process a data set unless it has a storage class assigned.

The SMSMANAGED parameter specifies whether or not the data set remains under DFSMS management. SMSMANAGED=NO clears the storage class for the data set and so forces the data set to be non-DFSMS-managed. SMSMANAGED=YES does not change the storage class; the data set keeps its original storage class assignment and remains DFSMS-managed.

Note: Setting SMSMANAGED to YES does not change a non-DFSMS-managed data set to DFSMS-managed because non-DFSMS-managed data sets are not processed by SMSACSTE.

Due to information that is available from OS/390 at the time this function is invoked, the only MAINVIEW SRM filter list parameters that have values are those that correspond to DFSMS ACS variables.

These parameters are:

DATACLAS	ENVIR	MGMTCLAS	SMSMANAGED
DD	EXPDT	NQUAL	STEPACCTn
DSN	HLQ	PGM	STOGROUP
DSNn	JOB	RACF	STORCLAS
DSNTYPE	JOBACCTn	RECORG	UNIT
DSORG	LLQ	RETPD	USER
DSTYPE	MAXSIZE	SIZE	VOL
			XMODE

ACS Variables

The ACS variables that are displayed by SMSACSTE shown in Table 3-7.

Table 3-7 ACS Variables Displayed by SMSACSTE

Variable	Explanation
ACCT_JOB ACCT_STEP ACSENVIR	First five job account fields First five step account fields ACS routine environment (RECALL, RECOVER, CONVERT, ALLOC, STORE, CHANGED, CTRANS, or other)
APPLIC DATA_CLASS DDN DSN DSNTYPE DSORG DSOWNER DSTYPE	Application (if RACF is installed) Data class DDname Data set name Data set name type Data set organization (PS, PO, VS, DA) Data set owner name (if RACF is installed) Data set type (GDS, PERM, TEMP). Note that the MAINVIEW SRM filter list parameter DSTYPE uses the value GDG rather than GDS, to be consistent with industry usage.
EXPDT GROUP JOB MANAGEMENT_CLASS MAXSIZE	Expiration date RACF group Job name Management class Size of primary extent plus all secondary extents
NQUAL NVOL PGM RECORG RETPD SIZE	Number of qualifiers in the data set name Number of volumes Program name Record organization (KS, ES, RR, LS) Retention period Size of primary extent or of primary + 1 secondary extent
STORAGE_CLASS STORAGE_GROUP UNIT USER VOL XMODE	Storage class Storage group Device type generic name User ID Volume serial number Execution mode (BATCH, TSO, STC)

Variables with a blank or null value are not displayed.

Messages

SMSACSTE issues message SVM0304I to display ACS variables:

SVM0304I *variable(value)*

where *variable* is the name of the ACS variable, and *value* is the value of that variable; for example:

```
SVM0304I  JOB(EMPCRMX) DDN(DD2) PGM(IEFBR14)
           DSN(EMPCRM.SMSACSTE.TEST) NQUAL(03)
SVM0304I  ACCT_JOB1(3500)
```

```

SVM0304I  SIZE(46KB) MAXSIZE(736KB) XMODE(JOB)
SVM0304I  DSNTYPE(NUL) DSTYPE(PERM) DSORG(PS)
SVM0304I  USER(EMPCRM) GROUP(BABEMP)
SVM0304I  ACSENVIR(ALLOC) UNIT(BABDA)
SVM0304I  STORAGE_CLASS(BBSC) STORAGE_GROUP(BBSG)

```

If SMSMANAGED=NO is specified, message SVM0307I SMSMANAGED=NO APPLIED is written to indicate that any storage class has been removed from the data set.

Note that the SMSACSTE function must be specified in the SMFUNCxx member to be available and must be set active to provide service.

Example

This example shows you how to use SMSACSTE to assist in debugging ACS routines.

SMFUNC04 member

```

SET  NAME=SMSACSTE ACTIVE=YES  The function is defined and
    MSG=I SMF=I                activated.
    FLST=04 RLST=02
    DESC='TEST ACS ROUTINES'

```

SMFLST04 member

```

SET  MODE=ACT                Select test data sets.
    INC DSN=SYSPGMR.ACSTEST./

```

SMRLST02 member

```

SET  SMSMANAGED=NO          If the third data set name qualifier
    INC DSN3=NOSMS*         begins with NOSMS, do not allow
                                the data set to be
                                DFSMS-managed.

SET  SMSMANAGED=YES         If the third data set name qualifier
    INC DSN3=SMS*           begins with SMS, let it stay
                                DFSMS-managed.

```

SMSMCREN—Identify Data Set Rename Requiring Management Class Change

The SMSMCREN function identifies data set renames that require a change in management class; that is, the management class for the new name is different from the management class for the old name.

Management classes allow enforcement of storage retention policies at a data set level. Management classes can affect data set retention/expiration date, release of allocated but unused space, and DFHSM migration and backup. When a data set is allocated, the management class can be set explicitly by JCL, or by an ACS routine. This management class stays with the data set until it is deleted, or manually changed by the data center through DFSMS utilities. Renaming the data set can reclassify that data set into a different management class; however, DFSMS does not change the management class on a data set name change.

For data set renames that require a management class change, SMSMCREN detects and reports the management class for the original data set name and the management class for the new data set name.

Rule List Parameters

SET Statement (in member SMRLSTxx)

Parameter	Description
EVENTID=xxxxx	Identifies a user event defined in an SMEVNTxx member

Rule List Parameter Explanations

EVENTID=

Purpose: Specifies the identifier assigned to a user event in SMEVNTxx. This parameter will cause an event to be generated from this function.

Note: If EVENTID= is used on an FLST SET statement with MODE=INACT, the event will still be issued.

Syntax: EVENTID=xxxxx

where xxxxx is the 5-character string specified on the EVNTID parameter in SMEVNTxx.

Default: None

Parameters Not Supported

The following filter list parameters are not supported for SMSMCREN:

ALCTYPE	BLKSIZE	BUFSP
CISIZE	CONTIG	CURSPACE
DATACLAS	DD	DEVTYPE
DIR	DSNTYPE	ERASE
EXPDT	EXTENT	IMBED
LRECL	MAXSIZE	MGMTCLAS
MNTYPE	OWNER	POOL
PRISPACE	RECFM	RELEASE
REPL	RETPD	REUSE
ROUND	SECSPACE	SIZE
SMS	STOGROUP	STORCLAS
UNIT	VOL	VSAMCOMP

Usage Notes

SMSMCREN operates on any SMS-managed data set being renamed. If the management class changes for the renamed data set, SMSMCREN issues a MAINVIEW SRM message. If the management class does not change with a rename, SMSMCREN does not issue a message.

SMSMCREN does not change the management class.

SMSMCREN recognizes a change in management class by comparing the existing management class with the management class returned by a call to the DFSMS ACS routine and the MAINVIEW SRM SMSACSMC function. If the management class function returns the same management class that is already assigned to the data set, SMSMCREN does not report a change; if the returned management class is different, SMSMCREN does report a change. Note that if the SMSACSMC function's rule that applies to this data set has REPLACE=NO, no management class is returned to SMSMCREN, which correctly interprets this as no change in management class. Also, if the SMSACSMC function's rule assigns a null management class (MGMTCLAS="), no change is reported by SMSMCREN (because this is identical to the output of REPLACE=NO).

All filter list parameters are available for use in the SMSMCREN filter list; however, the SMSACSMC function is limited to certain parameters. The function description of SMSACSMC in the EasySMS section of this manual details the filter list parameters available to SMSACSMC. In addition to those parameters, when SMSACSMC is called by SMSMCREN, the ENVIR parameter will be set to RENAM.

With the exception of the EVENTID rule list parameter, all resource selection must be done through the FLST.

Note that the SMSMCREN function must be specified in the SMFUNCxx member to be available and must be set active to provide services.

Messages

The MSG parameter for this function, either in SMFUNCxx or in SMFLSTxx, should not be set to W, E, S, or N, since this would prevent the informational message provided by SMSMCREN from being produced.

SMSMCREN generates message SMS3399I when it identifies a data set rename that requires a management class change:

```
SVM3399I  job,step,dd,dsn MGMTCLAS SHOULD BE ALTERED FROM
          oldclass TO newclass
```

where *oldclass* identifies the original management class and *newclass* identifies the new management class that should be assigned

```
SVM3400I  job,step,dsn SMSMCREN LOCATE FAILED RC = return code
SVM3401I  job,step,dsn SMSMCREN ACS CALL FAILED RC = return code
```

Example

This example shows you how to identify all renames that require a change in management class.

SMFUNCxx member

```
SET      NAME=SMSMCREN ACTIVE=YES
          MSG=I SMF=I
          FLST=76
          DESC='INVOKE SMS ON RENAME'
```

The function is defined and activated. There is no RLST specification.

SMFLST76 member

```
SET      MODE=ACT
          INC DSN=/'
```

This filter list specification is active. Selected resources are all data sets.

Example

This example shows you how to identify data set renames that require a change in management class, except for data sets that have a second name qualifier beginning with DEV or are being renamed by a job running in class T.

SMFUNCxx member

SET	NAME=SMSMCREN ACTIVE=YES	The function is defined and activated. There is no RLST specification.
	MSG=I SMF=I	
	FLST=MC	
	DESC='INVOKE SMS ON RENAME'	

SMFLSTMC member

SET	MODE=ACT	This filter list specification is active. It selects data sets without a second name qualifier beginning with DEV or in jobs in class T.
	EXC DSN2=DEV/	
	EXC JOBCLASS=T	

SMSSELECT—Control Selection of Volumes from SMS Storage Groups

The SMSSELECT function allows control over the selection of volumes from a storage group for DFSMS-managed data sets.

Volumes may be selected from the DFSMS-assigned storage group based on several criteria. SMSSELECT gives the data center the capability of assigning a specific volume based on the same volume selection techniques that are available in DASDPOOL:

- Best fit to a single extent on all volumes in the storage group.
- Volume with the largest percentage of free space.
- Volume with the largest single contiguous extent.
- Selection by performance statistics collected by MAINVIEW SRM.

DFSMS Subpools

SMS data sets can also be *subgrouped* to certain DFSMS volumes within a STORGRP. This is done by assigning SMSPOOLs to data sets with the parameters SMSPOOL and SMSPOOL_EXT. For these parameters to be processed, the system options SMS_ALLOC and SMS_EXTEND must be set to Yes. Please note that this functionality is limited to denying volumes outside of an DFSMS-defined pool. Volume selection cannot be directed outside of the data set's current STORGRP. If multiple pools are specified, the first pool with sufficient space for the data set is assigned. The first pool specified is the primary pool; all others are alternate pools. See Chapter 5, "How to Define DFSMS Subpools," for information on defining the SMSPOOL_{xx} definition parameter.

Critical Data Set Isolation

EasyPOOL allows you to identify a list of critical data sets in your environment that should never be placed on the same volume as any other data set in the list. EasyPOOL may also use data gathered by StorageGUARD about the EMC Symmetrix environments to ensure that critical data sets are not only on separate logical volumes but are placed on separate physical volumes to reduce data set contention.

Note: StorageGUARD must be installed and running to access EMC Symmetrix data.

It is important to consider, however, that CRITDSN is resource intensive and should be used only for a *small* list of critical data sets. It should not be used without considering the system impact. See Chapter 4, "How to Define Critical Data Set Lists," for information on defining the SMCRT_{xx} definition parameter.

System Parameters

The system parameters that affect SMSSELECT are described in Table 3-8.

Table 3-8 System Parameters—SMSSELECT

Parameter	Purpose
CRITLIST	Specifies the suffix of an SMCRITxx parameter member.
HISTDAYS	Specifies the number of days of data to retrieve from the StorageGUARD historical performance data collector and temporarily store in a data space. Values are 0-14 (7 or 14 recommended), which indicate the number of days of data to be retrieved. The default value is 0, which indicates that the data space is created, but no historical performance data is gathered and no HISTDPO pooling can be performed. If this value is changed, the historical performance data collector (SVSGP) must be stopped and restarted. For a valid HISTDPO test, do not specify a number of HISTDAYS higher than the total number of days collected in the historical performance data collector.
SMS_ALLOC	Specifies to EasyPOOL that SMSSELECT will be driven during DADSM ALLOCATE. If a POOL is coded in SMSPOOL, the current volume will be compared to the volumes in the POOL. If the current volume is not in a POOL assigned to the data set, the volume will be rejected with a DADSM return code of 4.
SMS_EXTEND	Specifies to EasyPOOL that SMSSELECT will be driven during DADSM EXTENDNV (new volume). If a POOL is coded in SMSPOOL_EXT, the current volume will be compared to the volumes in the POOL. If the current volume is not in a POOL assigned to the data set, the volume will be rejected with a DADSM return code of 4.

Rule List Parameters

SET Statement (in member SMRLSTxx)

Parameter	Description
CRITBIAS=nnnn	Interval for device selection based on performance
CRITEMC=YES/NO	Specifies whether the volume meeting CRITDSN criteria includes EMC physical volumes
CRITFAIL=YES/NO	Defines allocation if a volume meeting the criteria cannot be found
CRITLIST=xxxxxxx	Name of the table containing allocation volumes
DPORDEF=nnn	Default response time for device selection
DPORMAX=nnn	Maximum response time target for device selection
DPORMIN=nnn	Minimum response time target for device select
DPORSEP=nnn	Separation factor for device selection

Parameter	Description
DPOWIND= <i>nnnn</i>	Interval for device selection based on performance
EVENTID= <i>xxxxx</i>	Identifies a user event defined in an SMEVNT <i>xx</i> member
HDPORDEF= <i>nnn</i>	Default response time for device selection
HDPORMAX= <i>nnn</i>	Maximum response time target for device selection
HDPORMIN= <i>nnn</i>	Minimum response time target for device selection
HDPORSEP= <i>nnn</i>	Separation factor for device selection
HDPOSTIM= <i>nnnn</i>	Starting time each day for a range of snapshots to be analyzed
HDPOETIM= <i>nnnn</i>	Ending time each day for a range of snapshots to be analyzed
HDPODAYS=(<i>MO,TU,WE,TH FR,SA,SU</i>)	Day(s) of the week used in determining performance statistics for pooling
RAIDDEVTYPE= <i>xxxxxxxxx</i>	Allows the user to require a specific RAID device type for an allocation (EMC/RDFEMC/MIRROREMC/ PARITYEMC) (requires StorageGUARD) For RAIDDEVTYPE to be honored, a VOLSEL routine <i>must be</i> assigned.
SMSPOOL=(<i>xxxxxxxx,xxxxxxxx,...</i>)	Defines an SMSPOOL(s) to be used during DADSM ALLOCATE
SMSPOOL_EXT=(<i>xxxxxxxx,xxxxxxxx,...</i>)	Defines an SMSPOOL(s) to be used during DADSM EXTEND
VOLSEL= <i>xxxxxxxx</i>	Volume selection criteria (BESTFIT, PERCENT, MAXSPACE, DPO, HISTDPO, CRITDSN)

Rule List Parameter Explanations

CRITBIAS=

Purpose: Defines the number of data sets that can reside on the volume before the current allocation.

Syntax: CRITBIAS=*xxxx*

where *xxxx* is a number 1–9999.

Default: None

Note: This parameter works only in conjunction with VOLSEL=CRITDSN.

CRITEMC=

Purpose: Specifies whether the volume meeting CRITDSN criteria includes EMC physical volumes. If yes, data sets specified with CRITDSN should not reside on the same EMC Physical volume nor the same MVS Logical volume. If no, data sets specified with CRITDSN may reside on the same EMC physical volume. Default is no.

Syntax: CRITEMC=*YES/NO*

Default: CRITEMC=NO

Note: This parameter works only in conjunction with VOLSEL=CRITDSN.

CRITFAIL=

Purpose: Defines the allocation process if a volume meeting the criteria cannot be found. If Yes, allocation fails; if No, allocation is made to the best available volume. Like USELIM, if an acceptable volume cannot be found in the first pool, the best volume from the first pool is saved, and the system tries the next pool(s). If no acceptable volume is found, CRITFAIL is processed on the saved volume.

Syntax: CRITFAIL=*YES/NO*

Default: CRITFAIL=NO

Note: This parameter works only in conjunction with VOLSEL=CRITDSN.

CRITLIST=

Purpose: The table that contains allocation volumes.

Syntax: CRITLIST=*xxxxxxxx*

Default: None

Note: This parameter works only in conjunction with VOLSEL=CRITDSN.

DPORDEF=

Purpose: Specifies the default response time for device selection based on DASD performance. The default value is used when response time information has not been accumulated for a device, such as when it has just been brought online.

DPORDEF is useful only when specified in conjunction with DPORMIN/MAX.

Syntax: DPORDEF=*nnn*

where *nnn* specifies the default response time in milliseconds that will be substituted for unavailable information for a specific device.

Default: DPORDEF=50

DPORMAX=

Purpose: Specifies the maximum response time objective for device selection based on DASD performance.

DPORMAX is normally specified in conjunction with DPORMIN.

Note that the DASD Performance Optimization feature requires that VOLSEL=DPO be specified in order to select a device based on performance.

Syntax: DPORMAX=*nnn*

where *nnn* specifies the maximum response time in milliseconds that will be considered for device selection based on performance.

Default: None

DPORMIN=

Purpose: Specifies the minimum response time objective for device selection based on DASD performance.

DPORMIN is normally specified in conjunction with DPORMAX.

Note that the DASD Performance Optimization feature requires that VOLSEL=DPO be specified in order to select a device based on performance.

Syntax: DPORMIN=*nnn*

where *nnn* specifies the minimum response time in milliseconds that will be considered for device selection based on performance.

Default: None

DPORSEP=

Purpose: Specifies the response time value that is used to force selection of different volumes for data sets in the same jobstep. When multiple data sets are allocated in a single jobstep and device selection based on DASD performance is specified for some or all, this separation factor is added to the current response time for previously used volumes in order to increase the likelihood that new allocations will go to a different volume.

DPORSEP can only be specified in conjunction with DPORMIN/MAX.

Note that the DASD Performance Optimization feature requires that VOLSEL=DPO be specified in order to select a device based on performance.

Syntax: DPORSEP=*nnn*

where *nnn* specifies the response time in milliseconds that will be used to enhance data set separation across volumes

Default: DPORSEP=10

DPOWIND=

Purpose: Specifies the window (length of the performance interval) that is used to analyze the response characteristics of DASD devices for selection based on performance.

Note that the DASD Performance Optimization feature requires that VOLSEL=DPO be specified in order to select a device based on performance.

Syntax: DPOWIND=*nnnn*

where *nnnn* specifies the number of seconds over which the performance analysis will be made.

Default: DPOWIND=1800 (this is the maximum data retained in the performance tables)

EVENTID=

Purpose: Specifies the identifier assigned to a user event in SMEVNT_{xx}. This parameter will cause an event to be generated from this function.

Note: If EVENTID= is used on an FLST SET statement with MODE=INACT, the event will still be issued.

Syntax: EVENTID=*xxxxxx*

where *xxxxxx* is the 5-character string specified on the EVNTID parameter in SMEVNT_{xx}.

Default: None

HDPODAYS=

Purpose: Specifies the day(s) of the week to be used in determining performance statistics for pooling. This parameter is available only when VOLSEL=HISTDPO.

Syntax: HDPODAYS=(*MO, TU, WE, TH, FR, SA, SU*)

Default: HDPODAYS=(*MO, TU, WE, TH, FR, SA, SU*)

HDPOETIM=

Purpose: Specifies the ending time (each day) for a range of snapshots to be analyzed. This allows analysis of data from the same time range, or shift, each day. This parameter is available only when VOLSEL=HISTDPO.

Syntax: HDPOETIM=*nnnn*

where *nnnn* specifies a time in 24-hour format.

Default: HDPOETIM=2359

HDPORDEF=

Purpose: Specifies the default response time for device selection based on DASD performance. The default value is used when response time information has not been accumulated for a device, such as when it has just been brought online.

HDPORDEF is useful only when specified in conjunction with HDPORMIN/MAX. This parameter is available only when VOLSEL=HISTDPO.

Syntax: HDPORDEF=*nnn*

where *nnn* specifies the default response time in milliseconds that will be substituted for unavailable information for a specific device.

Default: HDPORDEF=50

HDPORMAX=

Purpose: Specifies the maximum response time objective for device selection based on DASD performance.

HDPORMAX is normally specified in conjunction with HDPORMIN. This parameter is available only when VOLSEL=HISTDPO.

Syntax: HDPORMAX=*nnn*

where *nnn* specifies the maximum response time in milliseconds that will be considered for device selection based on performance

Default: None

HDPORMIN=

Purpose: Specifies the minimum response time objective for device selection based on DASD performance.

HDPORMIN is normally specified in conjunction with HDPORMAX. This parameter is available only when VOLSEL=HISTDPO.

Syntax: HDPORMIN=*nn*

where *nn* specifies the minimum response time in milliseconds that will be considered for device selection based on performance.

Default: None

HDPORSEP=

Purpose: Specifies the response time value that is used to force selection of different volumes for data sets in the same jobstep. When multiple data sets are allocated in a single jobstep and device selection based on DASD performance is specified for some or all, this separation factor is added to the current response time for previously used volumes in order to increase the likelihood that new allocations will go to a different volume. This parameter is available only when VOLSEL=HISTDPO.

Syntax: HDPORSEP=*nn*

where *nn* specifies the response time in milliseconds that will be used to enhance data set separation across volumes.

Default: HDPORSEP=10

RAIDDEVTYPE=

Purpose: Allows the user to require a specific RAID device type for an allocation. If a particular hardware type is chosen, the candidate list of volumes will be limited to devices that have the requested attribute flag enabled. This parameter requires StorageGUARD to be installed and running.

Note: For this feature to be used under SMSSELECT, a data set must be assigned a VOLSEL parameter.

Syntax: RAIDDEVTYPE=xxxxxxxx

where xxxxxxxx is one of the following:

EMC The device must be an EMC RAID device.

RDFEMC The device must be EMC and have the RDF flag enabled.

MIRROREMC The device must be EMC and have the MIRROR flag enabled.

PARITYEMC The device must be EMC and have a parity flag enabled.

Default: None

SMSPOOL=

Purpose: Specifies 1 to 15 SMSPOOL(s) that can be used to limit volumes to be used in a primary allocation (the first volume) during DADSM ALLOCATE.

Syntax: SMSPOOL=(xxxxxxxx,xxxxxxxx,...)

where xxxxxxx is an SMSPOOL.

Default: None

Note: The SMSPOOL(s) must contain a subset of volumes from the SMS STORGRP assigned to the data set, or the allocation will fail. If multiple pools are specified, the first pool with sufficient space for the data set is assigned. The first pool specified is the primary pool; all others are alternate pools. In addition, SMS_ALLOC must be set to *Yes* under System parameters.

SMSPOOL_EXT=

Purpose: Specifies 1 to 15 SMSPOOL(s) that are to be used to limit volume selection during DADSM EXTENDNV on an extend to a new volume (anything after the first volume). Any SMS volumes can be coded in the pools. However, allocation to volumes will only be directed to volumes in the storage group assigned to the current data set.

Syntax: SMSPOOL_EXT=(xxxxxxxx,xxxxxxxx,...)

where xxxxxxxx is an SMSPOOL.

Default: None

Note: The SMSPOOL(s) must contain a subset of volumes from the SMS STORGRP assigned to the data set, or the allocation will fail. If multiple pools are specified, the first pool with sufficient space for the data set is assigned. The first pool specified is the primary pool; all others are alternate pools. In addition, SMS_EXTEND must be set to *Yes* under System parameters

VOLSEL=

Purpose: Specifies that volume selection within a pool should be based on current volume space availability.

Syntax: VOLSEL=BESTFIT/CRITDSN/DPO/HISTDPO/MAXSPACE/
PERCENT

Specifies that the volume with the smallest contiguous extent that satisfies the primary allocation should be selected.

CRITDSN Controls allocations by specifying data sets that should not reside on the same volume.

Warning! CRITDSN is resource intensive and should be used only for a *small* list of critical data sets. It should not be used without considering the system impact.

DPO Specifies volume selection based on performance statistics accumulated by MAINVIEW SRM.

HISTDPO Specifies volume selection based on historical volume performance.

Note: HISTDPO requires more resources than normal allocations.

MAXSPACE Specifies that the volume with the largest single contiguous extent should be selected.

PERCENT Specifies that the volume with the largest amount of free space should be selected.

Default: None

Parameters Not Supported

The following filter list parameters are not supported for SMSSELECT:

- GDGVER
- VALUE

Usage Notes

SMSSELECT provides services for DFSMS-managed data sets, both VSAM and non-VSAM, with a status disposition of NEW or MOD.

SMSSELECT does not require or use any information from the pool definitions in the SMPOOLxx parmlib member or the DASDPOOL filter/rule lists.

Tip: For RAIDDEVTYPE to be honored, a VOLSEL routine *must be* assigned.

Volume Selection

Volumes are selected from the storage group assigned to the DFSMS-managed data set. The VOLSEL rule list parameter specifies the method of volume selection. The selection techniques are described in Table 3-9.

Table 3-9 Volume Selection Techniques

Parameter	Explanation
BESTFIT	The volume with the smallest contiguous extent that satisfies the primary allocation amount is selected
PERCENT	The volume with the largest percentage of free space is selected
MAXSPACE	The volume with the largest single contiguous extent is selected
DPO	A volume is selected based on the performance statistics accumulated by the MAINVIEW SRM DASD Performance Optimization facility and the SMSSELECT rule list parameters DPORDEF, DPORMAX, DPORMIN, DPORSEP, and DPOWIND
CRITDSN	Controls allocations by specifying data sets that should not reside on the same volume. CRITDSN is resource intensive and should be used only for a <i>small</i> list of critical data sets. It should not be used without considering the system impact.

Messages

SMSSELECT issues message SVM3460I when a data set is assigned to a volume in a storage group:

SVM3460I *job,step,dd,dsn* VOL (*volser*) SELECTED

SMSSELECT issues one or more of the following messages:

SVM3353I *job,step,dd,dsn* CRITLIST = *crit_list*, CRITBIAS = *crit_bias*,
CRITFAIL = *crit_fail*, CRITEMC = *crit_emc*

SVM3354I *job,step,dd,dsn* CRITLIST FAILED ALL ELIGIBLE VOLUMES.
ALLOCATION REJECTED

SVM3360W *job,step,dd,dsn* ERROR ON CATALOG CALL. R15 = *rc*

SVM3361W *job,step,dd,dsn* ERROR ON RAID API CALL. R15 = *rc*

DASD Performance Optimization

The MAINVIEW SRM DPO facility is a specialized algorithm used during data set allocation to select a volume based on short-term historical performance data.

During I/O, the amount of time needed to complete the read or write operation is called the **response time**. The response time for a device can vary greatly due to the number of tasks competing for access to the volume and the channel path(s) to the volume. The DPO facility records device response times every 30 seconds and tables this performance data, maintaining 30 minutes worth of data in storage.

During allocation, the DPO facility scans the response time tables for the pooled devices over the indicated time interval; for example, for the last two minutes. The volume selected for allocation is the volume with the best response time. Alternatively, DPO can be used to select volumes based on response time objectives; for example, to select a volume that can provide a response time in the range of 20 to 40 milliseconds for some data sets or jobs and a volume that can provide a response time in the range of 40 to 60 milliseconds for other less response-critical data sets or jobs.

Within the MAINVIEW SRM ISPF interface, performance data can be displayed for each pool.

Starting and Stopping DPO

The DPO facility is started and stopped automatically as part of the SVOS started task.

Resources

DPO maintains a table of performance data above the 16MB line in extended private storage. The table contains one entry for each DASD UCB for 30 minutes (sampled twice a minute); each entry is 16 bytes long.

Usage

VOLSEL=DPO must be specified in the SMSSELECT rule for all data sets that are to use the DASD Performance Optimization allocation method.

The DPOWIND rule list parameter is used to specify the time interval over which the performance analysis is performed. For example, DPOWIND=120 indicates that the device with the best response time in the last 120 seconds should be selected for allocation. Since 30 minutes worth of data is retained, 1800 seconds is the logical upper limit for the DPOWIND operand.

Allocation based on response time objectives, rather than just the best response time, is specified in the DPORMIN and DPORMAX rule list parameters. These parameters set the minimum and maximum response times that are desired for the device to receive the allocation. For example, DPORMIN=20 and DPORMAX=40 will attempt to select a device that can provide a response time in the range of 20 to 40 milliseconds. If both DPORMIN and DPORMAX are zero, the best performing device is selected.

When multiple data sets are allocated in the same jobstep and some or all are using DPO allocation, the DPORSEP parameter is used to prevent all data sets from allocating to the same device, thus possibly degrading its performance significantly. The separation factor in DPORSEP is added to a previously used device's response time and the resulting figure is used in the selection process. This will artificially increase the device's response time, and cause subsequent allocations in the same job step to be allocated to another device with the desired response time characteristics.

In pools with fairly even performance statistics, the value of DPORSEP should generally be the average difference in response times on the different volumes. In pools with large differences in performance, such as pools that contain both cached and non-cached devices, allocations will concentrate on the fastest volumes, regardless of the separation factor.

In order to get a valid test of the DPO options, you should wait 30 minutes after you start DPO before you try to pool using DPO. It will take 30 minutes to collect the a full set of performance data. Any DPO-based allocations made before a full set of performance data has been collected may not reflect your actual selection algorithm.

Performance Data Display

The MAINVIEW SRM interface displays performance information for the devices in a pool. From the MPOOL view, type the **P** action line command to display performance data. The MVOLPER view (Figure 3-1) displays performance data (over the last 5 minute interval, maintained by DPO) for a selected pool, SMS pool, EMC physical device, or IBM physical device.

Figure 3-3 MVOLPER View

```

02JAN2001 13:29:57 ----- INFORMATION DISPLAY -----
COMMAND ==>
CURR WIN ==> 1          ALT WIN ==>
>W1 =MVOLPER=====SJSG=====*=====02JAN2001==13:29:57====MVSRM=====16
C Volume UCB      Resp Queue  Pend  Disc  Conn  Busy  Serv  EXCPs # of SSI
- ----- Addr   Time  Time   Time Time  Time  Time Time ----- UCBs Targ
  SMS011 85A3    4.43   0.00   0.17  0.00   4.26   0.03   4.26     3    1 SJSG
  SMS301 853E    4.43   0.00   0.21  0.00   4.22   0.03   4.22     3    1 SJSG
  SMS002 83BD    4.42   0.00   0.12  0.08   4.22   0.04   4.30     3    1 SJSG
  SMS303 8540    4.39   0.00   0.21  0.00   4.18   0.03   4.18     3    1 SJSG
  SMS001 8355    4.38   0.00   0.12  0.08   4.18   0.03   4.26     3    1 SJSG
  SMS004 8357    0.62   0.00   0.25  0.12   0.25   0.00   0.37     1    1 SJSG
  SMS21B 8524    0.50   0.00   0.25  0.00   0.25   0.00   0.25     1    1 SJSG
  SMS306 8543    0.50   0.00   0.25  0.00   0.25   0.00   0.25     1    1 SJSG
  SMS305 8542    0.50   0.00   0.25  0.00   0.25   0.00   0.25     1    1 SJSG
  SMS304 8541    0.50   0.00   0.25  0.00   0.25   0.00   0.25     1    1 SJSG
  SMS25A 8358    0.50   0.00   0.25  0.00   0.25   0.00   0.25     1    1 SJSG
  SMS003 8356    0.50   0.00   0.25  0.00   0.25   0.00   0.25     1    1 SJSG
  SMS302 853F    0.49   0.00   0.12  0.12   0.25   0.00   0.37     1    1 SJSG
  SMSG00 8576    0.37   0.00   0.12  0.00   0.25   0.00   0.25     1    1 SJSG
  SMSA00 85A2    0.00   0.00   0.00  0.00   0.00   0.00   0.00     0    1 SJSG
  SMSA01 8579    0.00   0.00   0.00  0.00   0.00   0.00   0.00     0    1 SJSG

```

Historical DASD Performance Optimization

The MAINVIEW SRM HISTDPO facility allows you to pool data sets based on performance data gathered over a period of up to two weeks. Historical performance data is obtained using historical performance data collector in the BMC Software StorageGUARD product, which must be licensed and operational at your site.

Data is stored in a data space to decrease the access time required to obtain the performance data during allocation. The data is averaged and compared against the options you choose to determine the optimum volume.

Usage

The system parameter, HISTDAYS, specifies the number of days of data to retrieve from the StorageGUARD historical performance data collector and temporarily store in a data space. Values are 0-14 (7 or 14 recommended), which indicate the number of days of data to be retrieved. The default value is 0, which indicates that the data space is created, but no historical performance data is gathered and no HISTDPO pooling can be performed. If this value is changed, the historical performance data collector (SVSGP) must be stopped and restarted. For a valid HISTDPO test, do not specify a number of HISTDAYS higher than the total number of days collected in the historical performance data collector.

VOLSEL=HISTDPO must be specified in the TAPEPOOL rule for all data sets that are to use the HISTDPO allocation method.

Starting time, ending time, and days of collection are specified in HDPOSTIM, HDPOETIM, and HDPODAYS respectively.

Allocation based on response time objectives, rather than just the best response time, is specified in the HDPORMIN and HDPORMAX rule list parameters. These parameters set the minimum and maximum response times that are desired for the device to receive the allocation. For example, HDPORMIN=20 and HDPORMAX=40 will attempt to select a device that can provide a response time in the range of 20 to 40 milliseconds. If both HDPORMIN and HDPORMAX are zero, the best performing device is selected.

When multiple data sets are allocated in the same jobstep and some or all are using HISTDPO allocation, the HDPORSEP parameter is used to prevent all data sets from allocating to the same device, thus possibly degrading its performance significantly. The separation factor in HDPORSEP is added to a previously used device's response time and the resulting figure is used in the selection process. This will artificially increase the device's response time, and cause subsequent allocations in the same job step to be allocated to another device with the desired response time characteristics.

In pools with fairly even performance statistics, the value of HDPORSEP should generally be the average difference in response times on the different volumes. In pools with large differences in performance, such as pools that contain both cached and non-cached devices, allocations will concentrate on the fastest volumes, regardless of the separation factor.

Example

This example shows you how to assign SMS volumes.

SMFUNC22 member

```
SET  NAME=SMSSELECT ACTIVE=YES
      MSG=I SMF=I
      FLST=53 RLST=26
      DESC='ASSIGN VOL FROM STOGROUP'
```

SMFLST53 member

```
SET  MODE=ACT                               Select all DFSMS data sets.
      INC STOGROUP=/
```

SMRLST26 member

```
SET  VOLSEL=BESTFIT                         Select the volume with the
      INC STOGROUP=APPLA                     extent that is closest to the
                                              data set's primary extent
                                              request for all data sets in
                                              storage group APPLA.
```

```
SET  VOLSEL=MAXSPACE                       Select the volume with the
      INC DSN=DIV325.PAY*.**.MSTR           largest single extent for payroll
                                              master files in division 325.
```

```
SET  VOLSEL=PERCENT                       Select the volume with the
      INC JOBCLASS=T                       largest percentage of free
                                              space for all DFSMS-managed
                                              data sets in jobs running in
                                              class T.
```

Example

This example shows you how to require all initial allocations (first volume) for an SMS-managed data set with a second qualifier of PARMS to allocate on SMS-managed volumes SMS001 to SMS009. This example requires that all extends to a new volume for SMS-managed data sets with a second qualifier of PARMS go to SMS-managed volumes SMS010 and SMS011.

Note: If the SMS-managed pools that are used in the SMSSELCT parms of SMSPool and SMSPool_EXT do not include volumes in the data sets STORGRP, the allocation will fail.

SMYS00 member

SMS_ALLOC=YES

Process SMSSELCT during DASD ALLOCATE processing.

SMS_EXTEND=YES

Process SMSSELCT during DADSM EXTENDNV processing.

SMFUNC22 member

SET NAME=SMSSELCT ACTIVE=YES
MSG=I SMF=Y
FLST=53RLST=26
DESC='ASSIGN VOL FROM STORGRP'

Define and activate the function.

SMSPOLxx member

SET NAME=SMSP1
INC VOL=SMS00/

Define SMSPOL named SMSP1 that includes all volumes that start with SMS00 (this will include SMS001 through SMS009).

SET NAME=SMSP2
INC VOL=(SMS010,SMS011)

Define SMSPOL named SMSP2 that includes all volumes SMS010 and SMS011).

SMFLST53 member

SET MODE=ACT
INC QUAL2=PARMS

Select all data sets that have PARMS as the second qualifier.

SET MODE=INACT
INC DADSM_FUNC=(ALLOCATE,EXTENDNV)

Do not process SMSPOL(s) for any other data sets. However, allow all other SMSSELCT options to be processed.

SMRLST26 member

SET SMSPOOL=SMSP1, SMSPOOL_EXT=SMSP2
INC QUAL2=PARMS

Allow the initial allocation only to the volumes defined in SMSP1, and allow extends to a new volume only to the volumes defined in SMSP2 only.

SPACCONV—Change Allocation Units to Blocks

The SPACCONV function converts space allocation values on new allocations from cylinders or tracks to blocks, for both primary and secondary allocation quantities.

Space allocation in blocks is advantageous because allocation in blocks is device independent; the system determines the number of tracks or cylinders necessary to hold the requested number of blocks, taking into account the device type. Allocation in units of cylinder or track is device-specific, because the size of the allocation unit varies with the device type.

SPACCONV is useful in allocating the correct amount of space when new device types are put into use, such as when converting from 3380 to 3390 devices. The old device type (for which the JCL space parameters are coded) is identified by the TRKLEN and TRKCYL parameters. SPACCONV automatically converts track and cylinder space requests into blocks, which, when applied by OS/390 during allocation to the new device type, allocates a proportionally larger or smaller total space.

System Parameters

The system parameters that affect SPACCONV are described in Table 3-10.

Table 3-10 **System Parameters—SPACCONV**

Parameter	Purpose
TRKCYL	Specifies the number of tracks per cylinder for the default device type. The value specified for 3380/3390/9345 devices should be fifteen. (Note that this specification is the same as the SCDS base configuration DEFINE under ISMF for DFSMS.) TRKCYL and TRKLEN are used by the DASDPOOL function to convert allocations in tracks or cylinders to megabytes for volume selection based on available space; for example, VOLSEL=BESTFIT. The information specified on these two parameters should reflect the devices that are most prevalent in your environment.
TRKLEN	Specifies the number of bytes per track for the default device type. Valid values are: 3380 – 47,476 3390 – 56,664 9345 – 46,456 Note that this specification is the same as the SCDS base configuration DEFINE under ISMF for DFSMS.) TRKCYL and TRKLEN are used by the DASDPOOL function to convert allocations in tracks or cylinders to megabytes for volume selection based on available space; for example, VOLSEL=BESTFIT. The information specified on these two parameters should reflect the devices that are most prevalent in your environment.

Rule List Parameters

SET Statement (in member SMRLSTxx)

Parameter	Description
EVENTID=xxxxx	Identifies a user event defined in an SMEVNTxx member
ROUND=YES/NO	Round allocation in blocks to cylinder boundary
TRKCYL=nn	Specifies the tracks per cylinder
TRKLEN=nnnnn	Specifies the bytes per track

Rule List Parameter Explanations

EVENTID=

Purpose: Specifies the identifier assigned to a user event in SMEVNTxx. This parameter will cause an event to be generated from this function.

Note: If EVENTID= is used on an FLST SET statement with MODE=INACT, the event will still be issued.

Syntax: EVENTID=xxxxx

where xxxxx is the 5-character string specified on the EVNTID parameter in SMEVNTxx.

Default: None

ROUND=

Purpose: Specifies that the converted space allocation should be rounded up to a cylinder boundary.

This is an optional parameter. If not used, the SMRLSTxx member can be omitted for this function.

Syntax: ROUND=YES/NO

Default: None

TRKCYL=

Purpose: Specifies the number of tracks per cylinder of the **source volume** for the SPACCONV function. The value is used to calculate the proper size of an allocation on new devices for specifications based on devices no longer used.

Syntax: TRKCYL=*nn*

where *nn* is a number in the range of 1–99.

Default: None (if not specified, the TRKCYL value in SMMSYSxx is used)

TRKLEN=

Purpose: Specifies the number of bytes per track of the **source volume** for the SPACCONV function. The value is used to calculate the proper size of an allocation on new devices for specifications based on devices no longer used.

Syntax: TRKLEN=*nnnnn*

where *nnnnn* is a number in the range of 1–99999.

Default: None (if not specified, the TRKLEN value in SMMSYSxx is used)

Parameters Not Supported

The following filter list parameters are not supported for SPACCONV:

- GDGVER
- VALUE

Usage Notes

SPACCONV applies to new DASD data sets (non-VSAM and JCL-defined VSAM). It does not apply to tape data sets or ISAM data sets. The original space allocation must be in cylinders or tracks; if it is already in blocks, no action is taken. Both primary and secondary allocation quantities are converted.

Conversion of space allocation from cylinders or tracks to blocks is based on the bytes/track and tracks/cylinder values specified in the TRKLEN and TRKCYL rule list parameters. The average block size is set to 9000 (this yields 5 blocks per track on both 3380 and 3390 devices). The calculated number of blocks is always increased by one. If ROUND is specified, rounding and alignment to a cylinder boundary is performed by OS/390. (Although an average block size of 9000 is used to calculate space, block size specified for the data set is not modified.)

If TRKLEN and/or TRKCYL are not specified in the rule list, the values are taken from the TRKLEN and TRKCYL parameters in the SMMSYSxx parmlib member. However, these SMMSYSxx parameters are used to represent the *normal* devices for allocation; accepting these values as a default causes the SPACCONV function to allocate the full size specified in the JCL.

The TRKLEN and TRKCYL specifications can have different values on different rules. This is useful when upgrading from two obsolete device types to a new device type, such as from 3350s and 3380s to 3390s.

Note: The ISPF data set information display in functions 3.4 and 3.2 do not always give the correct values for secondary space in blocks, particularly when the data set is allocated without a blocksize. The information provided by IEHLIST is correct, however.

Messages

The messages issued by SPACCONV are:

```
SVM3448I  job,step,dd,dsn SPACE CONVERTED TO prim/sec BLOCKS
SVM3318I  job,step,dd,dsn ROUND=y/n ASSIGNED
```

Note that the SPACCONV function must be specified in the SMFUNCxx member to be available and must be set active to provide services.

Example

This example shows you how to specify space conversion for unit specifications of 3350 and 3380.

SMFUNC33 member

SET	NAME=SPACCONV	ACTIVE=YES	The function is defined and
	MSG=I	SMF=I	activated.
	FLST=33	RLST=33	
	DESC='CONV SPACE TO BLOCKS'		

SMFLST33 member

SET	MODE=ACT	The function is to run in active mode
	INC UNIT=3350	for all allocation requests to devices
	INC UNIT=3380	no longer used (3350s and 3380s).

SMRLST33 member

SET	ROUND=YES	X	Convert space for allocations coded
	TRKLEN=47476	TRKCYL=15	for 3380 devices. Round allocations
	INC UNIT=3380		to an integral number of cylinders.
SET	TRKLEN=19069	TRKCYL=30	Convert space for allocations coded
	INC UNIT=3350		for 3350 devices.

SPACLIMI—Limit the Size of Space Allocation Requests

The SPACLIMI function limits the size of **new** data sets.

Setting and enforcing maximum sizes for data set space allocation is not directly possible in OS/390. Setting data set size limits is desirable because:

- Estimates coded in JCL SPACE statements may or may not be accurate.
- Excessive sizes for data sets which exist only for the life of a single job may never be detected, yet can cause other allocations to fail for lack of space.
- Unused space in permanent data sets may never be released (although the JCL parameter RLSE can be specified, there is no way to enforce this).

Excessive space allocations, both temporary and permanent, can unnecessarily consume significant DASD space, causing jobs to fail for lack of space and requiring the purchase of extra DASD devices.

The SPACLIMI function allows limits to be set on the size of the primary allocation for new data sets. Combined with the filter list parameter MAXSIZE, limits can also be placed on data sets based on their maximum possible size (primary plus the maximum secondary extents). Running the function in the simulate mode allows warnings to be given to users before tightly enforcing the limits.

System Parameters

The system parameters that affect SPACLIMI are described in Table 3-11.

Table 3-11 **System Parameters—SPACLIMI**

Parameter	Purpose
REJECT	Controls termination of processing at the first or last data set rejected by REJECT=YES in DASDPOOL or DSNCHECK. If termination is to take place on the first rejected data set, code REJECT=FIRST. If all data sets are to be processed before control is returned to OS/390 allocation, code REJECT=LAST.

Rule List Parameters

SET Statement (in member SMRLSTxx)

Parameter	Description
EVENTID=xxxxx	Identifies a user event defined in an SMEVNTxx member
LIMIT=nnnnnnnnKB/MB	Data set size limit in kilobytes or megabytes

Rule List Parameter Explanations

EVENTID=

Purpose: Specifies the identifier assigned to a user event in SMEVNTxx. This parameter will cause an event to be generated from this function.

Note: If EVENTID= is used on an FLST SET statement with MODE=INACT, the event will still be issued.

Syntax: EVENTID=xxxxx

where xxxxx is the 5-character string specified on the EVNTID parameter in SMEVNTxx.

Default: None

LIMIT=

Purpose: Specifies maximum allowable size for a data set primary allocation. This is a required parameter.

Syntax: LIMIT=nnnnnnnnKB/MB

where nnnnnnnn specifies the maximum allowed size of the new data set.

Default: None

Parameters Not Supported

The following filter list parameters are not supported for SPACLIMI:

- GDGVER
- VALUE

Usage Notes

SPACLIMI applies to *new* DASD data sets (VSAM and non-VSAM). If the requested amount of space for the data set being allocated is greater than the limit specified in the SPACLIMI rule, MAINVIEW SRM causes the job to fail with a JCL error. (A VSAM IDCAMS DEFINE will fail with a VSAM catalog return code of 68 – unable to allocate space on user volume.)

The limit is on the size of the primary allocation only. At allocation time, the requested primary allocation is converted to kilobytes using either the bytes/track and tracks/cylinder specifications in SMMSYSxx (TRKLEN and TRKCYL) or the average block length in the JCL SPACE parameter.

The LIMIT parameter must be specified; SPACLIMI takes no action without a rule specifying the limit for selected resources.

Note that the SPACLIMI function must be specified in the SMFUNCxx member to be available and must be set active to provide services.

The SPACRLSE function provides related space management functionality for new data sets.

Messages

The message issued by SPACLIMI when a limit is enforced is:

```
SVM3352I  job,step,dd,dsn REQUESTED SPACE nnnnKB/MB EXCEEDS LIMIT
```

where *nnnn* is the requested size in kilobytes or megabytes.

Example

This example shows you how to run SPACLIMI in simulate mode for production jobs and in active mode for test jobs and write SMF records only for production jobs.

SMFUNC1A member

```
SET  NAME=SPACLIMI ACTIVE=YES
      MSG=I SMF=N
      FLST=A1 RLST=A2
      DESC='LIMIT SPACE ALLOCATION'
```

The function is defined and activated. The default for SMF records is none.

SMFLSTA1 member

```
SET  MODE=SIM SMF=I MSG=N
      INC JOBCLASS=P
```

The function is to run in simulate mode for production jobs. All messages for SPACLIMI processing on production jobs are to be written to SMF for later analysis, but no messages are to be written to the production job logs.

```
SET  MODE=ACT
      INC JOBCLASS=T
```

The function is to run in active mode for test jobs, causing jobs to fail if a data set exceeds the limit.

SMRLSTA2 member

```
SET  LIMIT=50MB
      INC DSN=/'
```

Set the data set size limit for all data sets to 50 megabytes.

Example

This example shows you how to:

- Impose size limits on all data sets except those in pool OVERSZ and those in started tasks. The size limit for TSO data sets is 1 MB and 75 MB for batch jobs data sets not in STD0 pools.
- Restrict primary allocations for data sets in pools STD001-STD013 to 10% of their maximum size, with an overall limit of 100 MB.

SMFUNC20 member

```
SET NAME=SPACLIMI ACTIVE=YES
    MSG=I SMF=I
    FLST=39 RLST=3A
    DESC='ENFORCE SIZE LIMITS'
```

The function is defined and activated.

SMFLST39 member

```
SET MODE=ACT
    EXC POOL=OVERSZ
    EXC XMODE=STC
    INC DSN=/'
```

The function is to run in active mode. Data sets being allocated in pool OVERSZ or data sets being allocated from started tasks are exempted from size limits.

SMRLST3A member

```
SET LIMIT=1MB
    INC XMODE=TSO
```

For all TSO allocations, limit data sets to a 1 MB primary extent.

```
SET LIMIT=100MB
    INC MAXSIZE>1000MB POOL=STD0*
```

For STD0 pools, limit data sets potentially larger than 1000 MB to a 100 MB primary extent.

```
SET LIMIT=10MB
    INC MAXSIZE>100MB POOL=STD0*
```

For STD0 pools, limit data sets potentially larger than 100 MB to a 10 MB primary extent.

```
SET LIMIT=1MB
    INC MAXSIZE>10MB POOL=STD0*
```

For STD0 pools, limit data sets potentially larger than 10 MB to a 1 MB primary extent.

```
SET LIMIT=100KB
    INC MAXSIZE>1MB POOL=STD0*
```

For STD0 pools, limit data sets potentially larger than 1 MB to a 100 KB primary extent.

```
SET LIMIT=50KB
    INC DSN=/' POOL=STD0*
```

For STD0 pools, limit data sets with a maximum size less than 1 MB to a 50 KB primary extent.

```
SET LIMIT=75MB
    INC XMODE=JOB
```

For all batch job allocations except those in the STD0 pools, limit data sets to 75 MB.

SPACRLSE—Release Data Set Space at Close

The SPACRLSE function causes unused space in physical-sequential data sets to be released when the data set is closed.

Inaccurate space estimates can allocate significantly more space in data sets than is actually needed. Excessive space requests are frequently made when the expected data set size is not known, particularly to avoid abending with a space related problem. While OS/390 provides a facility to release unused space from a new data set, it must be requested using the RLSE subparameter of the SPACE parameter in JCL. If RLSE is not specified in the JCL, the extra space in the data set is not freed until the data set is eventually deleted, or until the storage administrator manually searches for and releases that space.

Excessive space allocations unnecessarily consume significant DASD space, causing other jobs to fail for lack of space and requiring the purchase of extra DASD devices.

The SPACRLSE function allows automated release of unused space, regardless of JCL specifications.

Rule List Parameters

SET Statement (in member SMRLSTxx)

Parameter	Description
EVENTID=xxxxx	Identifies a user event defined in an SMEVNTxx member
RLSE=ALL/SEC/NO	Specifies whether space is released for <i>all</i> selected data sets, for only those selected data sets that have <i>secondary</i> allocations, or <i>no</i> data sets.

Rule List Parameter Explanations

EVENTID=

Purpose: Specifies the identifier assigned to a user event in SMEVNTxx. This parameter will cause an event to be generated from this function.

Note: If EVENTID= is used on an FLST SET statement with MODE=INACT, the event will still be issued.

Syntax: EVENTID=xxxxx

where xxxxx is the 5-character string specified on the EVNTID parameter in SMEVNTxx.

Default: None

RLSE=

Purpose: Specifies whether space is released for *all* selected data sets, for only those selected data sets that have *secondary* allocations, or *no* data sets. This is a required parameter.

Syntax: RLSE=ALL/SEC/NO

ALL	releases space for all data sets
SEC	releases space for data sets with a secondary allocation
NO	turns the release flag off

Default: None

Parameters Not Supported

The following filter list parameters are not supported for SPACRLSE:

- GDGVER
- VALUE

Usage Notes

SPACRLSE provides services for DASD physical sequential data sets. Unused space is released from the data set when the data set is closed. SPACRLSE does not release the space itself, but sets the RLSE JCL parameter for selected resources.

SPACRLSE works for both new allocations and existing data sets.

Note: SPACRLSE does not work on data sets allocated through ISPF service 3.2 (data set utilities).

The RLSE parameter specifies the scope of the function. *ALL* specifies that all selected data sets should be flagged for space release. *SEC* specifies that only data sets with a secondary allocation should be flagged for release; data sets with only a primary allocation are ignored. This allows the primary allocation to remain at full size, which may be desirable if the data set is extended at a later time.

The RLSE parameter must be specified; SPACRLSE takes no action without a rule specifying the release scope for selected resources.

SPACRLSE does not modify the action of OS/390 in processing the release of unused data set space; therefore, OS/390 restrictions still apply. See the discussion of the RLSE subparameter of the SPACE parameter in the *JCL Reference Manual* for your release of OS/390.

Note that the SPACRLSE function must be specified in the SMFUNCxx member to be available and must be set active to provide services.

Messages

If RLSE=SEC or ALL, the message issued by SPACRLSE is:

SVM3349I *job,step,dd,dsn* RELEASE PARAMETER ASSIGNED

If RLSE=NO, the message issued by SPACRLSE is:

SVM3349I *job,step,dd,dsn* RELEASE PARAMETER REMOVED

Example

This example shows you how to release space for all data sets except certain temporary work files which are frequently extended during the course of a job; for production jobs, do not release space in data sets with only a primary allocation.

SMFUNCD5 member

```
SET  NAME=SPACRLSE           The function is defined and activated.
      ACTIVE=YES
      MSG=I SMF=I
      FLST=D3 RLST=D3
      DESC='FREE UNUSED SPACE'
```

SMVARSD0 member

```
SET  VARIABLE=DLYWORK       Define a variable which identifies all
      VALUE=**.DLYWORK*      transitory work files that should not
      VALUE=**.TRAN*/        have free space deleted.
      VALUE=**.TEMP*
      VALUE=WORK/
```

SMFLSTD3 member

```
SET  MODE=ACT               Select all data sets for SPACRLSE
      EXC DSN=&DLYWORK       processing, except those specified in
      INC DSN=/              the variable &DLYWORK.
```

SMRLSTD3 member

```
SET  RLSE=SEC               For production jobs, release space only
      INC JOBCLASS=P         if the data set has a secondary
                              allocation.

SET  RLSE=ALL               For all other data sets, release space.
      INC DSN=/'
```

SPACSQTY—Set Primary and Secondary Space for Data Set

The SPACSQTY function sets primary and secondary space allocation values for non-VSAM and VSAM DASD data sets being created. Exceptions for VSAM allocations are noted in the parameters explanations.

Primary and secondary space allocation values have a significant role in effective use of DASD space. Allocations which are consistently too high, even for temporary data sets, utilize more space than is necessary, possibly preventing other jobs from allocating needed resources. Allocations which are too low require subsequent processing of some type, either automated or manual, to allocate more space to the data set.

SPACSQTY allows the data center to set requirements and limitations on data set size and to enforce those requirements rigorously. Application of predefined primary and secondary space allocations can be based on data set names, on device types, on assigned pools, on size or maximum possible size, or on any property allowed as a MAINVIEW SRM filter list parameter.

Rule List Parameters

SET Statement (in member SMRLSTxx)

Parameter	Description
ALCTYPE=xxx	Quantity unit of space allocation (TRK, CYL, KAV, MAV, UAV, BLK, KB, MB) that applies to both primary and secondary space quantities
AVL=nnnnn	Average block or record length
CONTIG=YES/NO	Specifies whether a data set is allocated with contiguous space required
DIR=nnnn	Sets number of directory blocks for partitioned data sets
EVENTID=xxxxx	Identifies a user event defined in an SMEVNTxx member
PQTY=nnnnn	Primary space allocation in unit specified
REPLACE=YES/NO	Allows override to JCL-specified space allocations
ROUND=YES/NO	Rounds space allocated in blocks to nearest cylinder
SQTY=nnnnn	Secondary space allocation in unit specified

Rule List Parameter Explanations**ALCTYPE=**

Purpose Specifies the unit of space allocation. Applies to both primary and secondary space quantities.

Syntax: ALCTYPE=*xxx*

where *xxx* is a valid allocation type from the following list:

KB A block allocation where PQTY and SQTY are assumed to be kilobytes

MB A block allocation where PQTY and SQTY are assumed to be megabytes

TRK Tracks

CYL Cylinders

KAV AVGREC=K

MAV AVGREC=M

UAV AVGREC=U

BLK Blocks

Note: KAV, MAV, and UAV are not valid parameters for VSAM allocations.

Default: KB

AVL=

Purpose Specifies the average block size for BLK allocations and record length for KAV, MAV, and UAV allocations.

Syntax: AVL=*nnnnn*

where *nnnnn* is a number in the range 1 – 32,767.

Default: 9000

CONTIG=

Purpose: Specifies whether a data set is allocated with contiguous space required.

Syntax: CONTIG=*YES/NO*

Default: None

DIR=

Purpose Specifies the number of directory blocks for partitioned data sets.

Note: This parameter can only be used to change the directory blocks. The user must request DSORG=PO or directory blocks in the original allocation.

Syntax: DIR=*nnnn*

where *nnnn* is a number between 1 and 4096.

Default: None

EVENTID=

Purpose: Specifies the identifier assigned to a user event in SMEVNT*xx*. This parameter will cause an event to be generated from this function.

Note: If EVENTID= is used on an FLST SET statement with MODE=INACT, the event will still be issued.

Syntax: EVENTID=*xxxxx*

where *xxxxx* is the 5-character string specified on the EVNTID parameter in SMEVNT*xx*.

Default: None

PQTY=

Purpose: Specifies the primary space allocation quantity in the unit specification ALCTYPE for a new non-VSAM data set.

Syntax: PQTY=*nnnnn*

where *nnnnn* is a number in the range 1–99999.

Default: None

REPLACE=

Purpose: Specifies whether the space allocations already specified for the data set (in a JCL, for example) are overridden. If REPLACE=NO, and a primary or a secondary space allocation already exists for the new data set, then the SPACSQTY function does not replace either space allocation value. If REPLACE=YES, the SPACSQTY function sets the primary and secondary space allocations, regardless of any existing space parameters.

Syntax: REPLACE=*YES/NO*

Default: REPLACE=*NO*

Note: REPLACE only applies to PQTY and SQTY parameters. CONTIG and ROUND are not affected by REPLACE.

ROUND=

Purpose: Specifies that the space allocation be rounded to the nearest cylinder for new data sets.

Note that this rounding process is performed by OS/390; MAINVIEW SRM merely switches on the appropriate bit to allow OS/390 to perform this service.

Syntax: ROUND=*YES/NO*

Default: None

Note: Both SPACCONV and SPACSQTY functions allow the ROUND parameter to be specified. If both are specified for the same data set, the value in SPACCONV will be the value assigned to the data set, because SPACCONV receives control after SPACSQTY.

SQTY=

Purpose: Specifies the secondary space allocation quantity in the units specified in ALCTYPE for a new non-VSAM data set.

Syntax: SQTY=nnnnn

where *nnnnn* is a number in the range 1–99999.

Default: None

Parameters Not Supported

The following filter list parameters are not supported for SPACSQTY:

- GDGVER
- VALUE

Usage Notes

SPACSQTY sets the data set's primary and secondary space allocation in the units specified in the ALCTYPE parameter.

Note: For compatibility with earlier MAINVIEW SRM releases, PQTY and SQTY are assumed to be in kilobytes if ALCTYPE is not coded. The resulting allocation will be calculated in average blocks based on the AVL parameter.

SPACSQTY provides services for VSAM and non-VSAM DASD data sets being created; old data sets are not affected by SPACSQTY

Note: ALCTYPE applies to both primary and secondary space quantities. If, for example, PQTY and ALCTYPE are set for an allocation, ALCTYPE also applies to the secondary that was originally assigned to the allocation.

The control parameters (ALCTYPE, AVL, CONTIG, DIR, ROUND) allow full control over space allocation specifications. ALCTYPE specifies the unit of space that is referred to by PQTY and SQTY; for example, PQTY=10 ALCTYPE=CYL requests 10 cylinders of space.

The SPACSECA function also allows specification of secondary allocation quantity. SPACSECA, however, is invoked when a data set is out-of-space on its primary extent.

Note: The ISPF data set information display in functions 3.4 and 3.2 do not always give the correct values for secondary space in blocks, particularly when the data set is allocated without a blocksize. The information provided by IEHLIST is correct, however. Note that the SPACSQTY function must be specified in the SMFUNCCxx member to be available and must be set active to provide services.

REPLACE parameter

The REPLACE parameter allows SPACSQTY to be applied in either a restrictive or a lenient manner. The REPLACE parameter specifies whether a previously-specified space allocation is to be overridden by the SPACSQTY PQTY and SQTY parameters.

- REPLACE=YES forces the use of the space allocation established in SPACSQTY, regardless of any space parameters already specified in the JCL or elsewhere.
- REPLACE=NO accepts the original space parameters. The space allocation set by SPACSQTY is applied only if there is no existing space allocation specification.

Messages

SPACSQTY generates message SVM3446I when it provides new space allocation values for a data set:

```
SVM3446I  job,step,dd,dsn SPACE SET TO: TYPE=alloc_type,
          PRIMARY=prim_alloc, SECONDARY=sec_alloc, DIRECTORY
          BLOCKS=dir_blocks, AVL=avl_length
```

Messages issued by SPACSQTY include:

```
SVM3317I job,step,dd,dsn CONTIG y/n ASSIGNED
SVM3318I job,step,dd,dsn ROUND y/n ASSIGNED
```

Controlling Extent Size

In general, extent-size during allocation/volume-addition can be controlled by several StopX37/II functions when at least one of the eligible volumes contain an extent large enough to satisfy the requested size or when SPACSWIR is active, and none of the eligible volumes contain an extent large enough to satisfy the requested size. The StopX37/II functions are:

- SPACSQTY adds or overrides requested space on new allocations.
- SPACPRIM reduces the size of the primary extent if the requested value cannot be found.
- SPACSWIR reduces the size of the first allocation on an added volume if the requested value cannot be found.
- SPACVOLA adds a volume when OS/390 encounters an end-of-volume condition:

Example

This example shows you how to force a primary and secondary space allocation limit on all new data sets being allocated on volumes starting with P and whose data set name ends with TEST, and with a normal termination disposition of DELETE.

SMFUNCxx member

```
SET  NAME=SPACSQTY ACTIVE=YES
      MSG=I SMF=E
      FLST=JV RLST=JV
      DESC='SET PRIM AND SEC SPACE QTY'
```

The function is defined and activated.

SMFLSTJV member

```
SET  MODE=ACT
      INC LLQ=TEST/ VOL=P/
```

This filter list specification is active. It selects all data sets whose final qualifier begins with TEST, and which are being allocated to any volume whose name begins with P.

SMRLSTJV member

```
SET  PQTY=100 SQTY=20KB REPLACE=YES  X
      ALCTYPE=KB
      INC DISP2=DELETE
```

All data sets selected by the FLST specification above and with a normal termination disposition of DELETE are given a primary and secondary space allocation as shown. Any space allocation parameters in the JCL are overridden.

Example

This example shows you how to:

- Impose space limitations determined by the fourth account parameter of the job step EXEC card for all data sets going on a public work device (POOL=PUBWRK), except for those jobs with a third job account parameter beginning with PR.
- Apply this limitation in simulation mode only, to determine impact.

SMFUNCxx member

```
SET  NAME=SPACSQTY ACTIVE=YES
      MSG=I SMF=I
      FLST=23 RLST=22
      DESC='SET PRIM AND SEC SPACE QTY'
```

The function is defined and activated.

SMFLST23 member

```
MODE=SIM
INC POOL=PUBWRK
EXC JOBACCT3=PR/
```

This filter list specification is in simulation mode only. It selects all data sets in pool PUBWRK.

SMRLST22 member

```
SET  PQTY=1000 SQTY=500 REPLACE=YES ALCTYPE=KB
      INC STEPACCT3=LRG
SET  PQTY=500 SQTY=100 REPLACE=YES ALCTYPE=KB
      INC STEPACCT3=MOD
SET  PQTY=100 SQTY=25 REPLACE=YES ALCTYPE=KB
      INC STEPACCT3=SML
SET  PQTY=50 ALCTYPE=CYL CONTIG=YES
      REPLACE=YES
      INC DSN=ACCT/ DSORG=PS SIZE<5MB
```

Force a primary and secondary space allocation on data sets selected by the filter list above, depending on the value in the third account field of the step EXEC statement.

Allocate 50 cylinders for all PS accounting data sets that are allocated with less than 50 MB.

STKSUPP—Split Unit Affinity Specifications Over Devices

The STKSUPP function splits unit affinity specifications across separate devices to eliminate requirements to enter or remove volumes from silos. STKSUPP also allows specification of default units for scratch volumes outside of silos.

When unit affinity JCL parameters force multiple DDs onto a single drive, volumes may be required to be entered into or removed from silos. HSCs software component determines the location of the volume specified in the *primary* DD statement; volumes on *secondary* DD statements have to be moved to use the unit allocated to the affinity group.

STKSUPP alleviates the problem by splitting the affinity into groups based on tape location—for example, into a cartridge drive group and a silo group. While the affinity split requires more drives to run the job, it does not require operator intervention to enter or remove volumes from silos.

System Parameters

The system parameter that affects STKSUPP is described in Table 3-12.

Table 3-12 **System Parameters—STKSUPP**

Parameter	Purpose
STKSCR	Specifies the default location of scratch tapes for the STKSUPP function.

Rule List Parameters

SET Statement (in member SMRLST_{xx})

Parameter	Description
DEFUNIT=xxxxxxx	Generic unit name for volumes outside the silos
EVENTID=xxxxx	Identifies a user event defined in an SMEVNT _{xx} member
SPLIT= <u>YES</u> /NO	Specifies whether to split unit affinities with STK silos

Rule List Parameter Explanations**DEFUNIT=**

Purpose: Specifies the default unit name for volumes located outside STK silos. Under some circumstances with silos in separate physical locations, the HSC processing can cause a non-silo tape to be allocated to a silo device in a remote location, possibly without an operator. The DEFUNIT parameter forces non-silo volumes to be allocated to a local silo.

Syntax: DEFUNIT=xxxxxxx
where xxxxxxx specifies a generic unit name.

Default: DEFUNIT=3480

EVENTID=

Purpose: Specifies the identifier assigned to a user event in SMEVNTxx. This parameter will cause an event to be generated from this function.

Note: If EVENTID= is used on an FLST SET statement with MODE=INACT, the event will still be issued.

Syntax: EVENTID=xxxxx
where xxxxx is the 5-character string specified on the EVNTID parameter in SMEVNTxx.

Default: None

SPLIT=

Purpose: Specifies whether affinity separation should be applied to resources selected for the STKSUPP function. When unit affinity directs multiple DD statements to a single drive, STKSUPP can apply affinity separation to allocate separate devices for silo and non-silo volumes, so volumes would not have to entered into or removed from silos.

Syntax: SPLIT=YES/NO

Default: SPLIT=YES

Parameters Not Supported

The following filter list parameters are not supported for STKSUPP:

ALCTYPE	BUFSP	CATALOG
CONTIG	CURSPACE	DSORG
ENVIR	ERASE	EXTENT
GDGVER	IMBED	MAXSIZE
OWNER	PRISPACE	RECORG
REPL	REUSE	ROUND
SECSPACE	SIZE	VALUE
VSAMCOMP	VSAMDEF	

Usage Notes

When unit affinity is used to group DD statements onto a single drive in conjunction with StorageTek tape silos, the first *primary* DD statement's drive is allocated for the group, and volumes on subsequent DD statements are required to use that drive. This can require tapes to be entered into or removed from tape silos. For example:

```
//STEP10 EXEC  PGM=....  
//DD1      DD   DSN=DSN1,UNIT=CART,VOL=SER=SILOA1,DISP=OLD  
//DD2      DD   DSN=DSN2,UNIT=AFF=DD1,VOL=SER=TAPE01,DISP=OLD  
//DD3      DD   DSN=DSN3,UNIT=AFF=DD1,VOL=SER=SILOB1,DISP=OLD
```

HSC requires OS/390 to select silo A (for tape SILOA1); the affinities on DD2 and DD3 then require the non-silo tape TAPE01 and the silo B tape SILOB1 to be entered into silo A.

The STKSUPP function can be directed to split the affinities so that tapes are not required to be relocated in or out of silos. For example:

```
SET  SPLIT=YES  
INC  UNIT=CART
```

With this specification in the rule list for STKSUPP, the above JCL would require 3 devices: silo A for DD1, a cartridge drive for DD2, and silo B for DD3. This is the equivalent of:

```
//STEP10 EXEC PGM=....
//DD1 DD DSN=DSN1,UNIT=SILOA,VOL=SER=SILOA1,DISP=OLD
//DD2 DD DSN=DSN2,UNIT=CART,VOL=SER=TAPE01,DISP=OLD
//DD3 DD DSN=DSN3,UNIT=SILOB,VOL=SER=SILOB1,DISP=OLD
```

STKSUPP does not remove unit affinity from devices within the same silo or from *manual* drive groups. For example:

```
//STEP10 EXEC PGM=....
//DD1 DD DSN=DSN1,UNIT=CART,VOL=SER=SILOA1,DISP=OLD
//DD2 DD DSN=DSN2,UNIT=AFF=DD1,VOL=SER=TAPE01,DISP=OLD
//DD3 DD DSN=DSN3,UNIT=AFF=DD2,VOL=SER=SILOB1,DISP=OLD
//DD4 DD DSN=DSN4,UNIT=AFF=DD3,VOL=SER=TAPE02,DISP=OLD
//DD5 DD DSN=DSN5,UNIT=AFF=DD4,VOL=SER=SILOB2,DISP=OLD
//DD6 DD DSN=DSN6,UNIT=AFF=DD5,VOL=SER=SILOA2,DISP=OLD
```

Without affinity separation processing by STKSUPP, tapes TAPE01 and TAPE02 would have to be entered into silo A, and tape SILOB1 and SILOB2 would have to be ejected from silo B and entered into silo A (tapes SILOA1 and SILOA2 are already in silo A). STKSUPP separates affinity groups for silo A, silo B, and a cartridge drive, as if the following JCL had been coded:

```
//STEP10 EXEC PGM=....
//DD1 DD DSN=DSN1,UNIT=SILOA,VOL=SER=SILOA1,DISP=OLD
//DD2 DD DSN=DSN2,UNIT=CART,VOL=SER=TAPE01,DISP=OLD
//DD3 DD DSN=DSN3,UNIT=SILOB,VOL=SER=SILOB1,DISP=OLD
//DD4 DD DSN=DSN4,UNIT=AFF=DD2,VOL=SER=TAPE02,DISP=OLD
//DD5 DD DSN=DSN5,UNIT=AFF=DD3,VOL=SER=SILOB2,DISP=OLD
//DD6 DD DSN=DSN6,UNIT=AFF=DD1,VOL=SER=SILOA2,DISP=OLD
```

Scratch Tapes

STKSUPP cannot determine the appropriate location for scratch tapes. The STKSCR system parameters specified in the SMMSYSxx parmlib member define default locations for the following scratch requests: standard-label tapes, non-labeled tapes, ASCII tapes, and non-standard-label tapes.

Each type of scratch request can be directed to any silo (operand IN), a specific silo ID (*nnn*), or to devices outside of silos (operand OUT). For example:

```
STKSCR=( IN , OUT , OUT , OUT )
```

directs standard-label scratch tapes to silo devices and all other scratch tapes to non-silo devices.

In the following job, without STKSUPP and the preceding STKSCR parameters, all tapes would have to be moved to the device associated with DD1.

```
//STEP10 EXEC PGM=...
//DD1 DD DSN=DSN1,UNIT=CART,DISP=NEW,LABEL=(1,SL)
//DD2 DD DSN=DSN2,UNIT=AFF=DD1,DISP=NEW,LABEL=(1,NL)
//DD3 DD DSN=DSN3,UNIT=AFF=DD2,DISP=NEW,LABEL=(1,NL)
//DD4 DD DSN=DSN4,UNIT=AFF=DD3,DISP=NEW,LABEL=(1,SL)
//DD5 DD DSN=DSN5,UNIT=AFF=DD4,VOL=SER=SILOB1,DISP=OLD
```

With the STKSCR parameters above, affinity separation splits the scratch requests into two groups: all standard-label scratch tapes and all other scratch tapes. It also splits out DD5, since it resides in silo B. This is equivalent to the following JCL:

```
//STEP10 EXEC PGM=...
//DD1 DD DSN=DSN1,UNIT=SILOx,DISP=NEW,LABEL=(1,SL)
//DD2 DD DSN=DSN2,UNIT=CART,DISP=NEW,LABEL=(1,NL)
//DD3 DD DSN=DSN3,UNIT=AFF=DD2,DISP=NEW,LABEL=(1,NL)
//DD4 DD DSN=DSN4,UNIT=AFF=DD1,DISP=NEW,LABEL=(1,SL)
//DD5 DD DSN=DSN5,UNIT=AFF=SILOB,VOL=SER=SILOB1,DISP=OLD
```

Default Unit Names

Non-silo tapes allocated to silo devices can be allocated to remote silos, if your data center has silos that are physically separate. When OS/390 attempts to allocate a cartridge device and the HSC software changes the allocation to a silo, the original unit satisfies the request. OS/390 may then select a device in a remote silo, which may not have an operator.

The DEFUNIT rule list parameter on the STKSUPP rule list can override the unit specified for a non-silo tape request, ensuring that cartridges allocated to a silo are allocated to a local device. For example:

```
SET SPLIT=YES DEFUNIT=SILOB
INC UNIT=CART
INC UNIT=3480
```

With this specification, any non-silo tape request with CART or 3480 specified in the JCL will be allocated to silo B (any device represented by the generic name SILOB).

Messages

STKSUPP issues message SVM3600I when affinity processing for the specified ddname is changed.

SVM3600I AFFINITY MODIFIED ON *ddname* + *nnnnn* FROM *nnnn(nn)* TO *nnn*

STKSUPP issues message SVM3601I when unable to split the affinity due to a shortage of devices.

SVM3601W AFFINITY SPLIT ON *ddname* WOULD EXCEED AVAILABLE DEVICES

Example

This example shows you how to split affinities to prevent tape movement in and out of silos.

SMMSYS00 member

SET STKSCR=(IN,OUT,OUT,OUT)

Standard-label scratch tapes are allocated inside silos; all other scratch tapes are allocated outside silos.

SMFUNC00 member

SET FUNC=STKSUPP ACTIVE=YES
MSG=I SMF=N
FLST=25 RLST=25
DESC='STK UNIT AFFINITY SPLIT'

The function is active.

SMFLST25 member

SET MODE=ACT
INC UNIT=TAPE
INC UNIT=CART
INC UNIT=SILO*

All allocations coded for TAPE, CART, or any unit starting with SILO are selected for processing.

SMRLST25 member

SET SPLIT=NO
INC JOB=PAYR* JOBCLASS=V

Do not split affinities for payroll jobs in class V.

SET SPLIT=YESDEFUNIT=CART
INC UNIT=TAPE
INC UNIT=CART
INC UNIT=SILO*

Otherwise, separate affinities for allocations requesting TAPE, CART, or SILO units.

SUPJSCAT—Suppress JOBCAT and STEPCAT DD Statements

The SUPJSCAT function suppresses the action of all JOBCAT and STEPCAT DD statements in a job.

JOBCAT and STEPCAT statements allow circumvention of the cataloging structure implemented by the data center. Jobcat/stepcats allow data sets to be cataloged in a catalog that does not have an entry in the OS/390 master catalog. These data sets are then accessible only through a jobcat/stepcat in the JCL (or a hard-coded volume ID); they cannot be located by the normal OS/390 catalog search process. This can create problems if access is needed to the data sets by personnel who are not aware of the jobcat/stepcat requirements.

JOBCATs and STEPCATs are not supported by DFSMS; use of these DD statements in an DFSMS-managed system causes a JCL error.

SUPJSCAT provides a method of monitoring jobs for jobcat/stepcat usage without causing problems (in simulation mode) and makes JCL changes unnecessary for the removal of jobcat/stepcat cards.

Rule List Parameters

SET Statement (in member SMRLSTxx)

Parameter	Description
EVENTID=xxxxx	Identifies a user event defined in an SMEVNTxx member

Rule List Parameter Explanations

EVENTID=

Purpose: Specifies the identifier assigned to a user event in SMEVNTxx. This parameter will cause an event to be generated from this function.

Note: If EVENTID= is used on an FLST SET statement with MODE=INACT, the event will still be issued.

Syntax: EVENTID=xxxxx

where xxxxx is the 5-character string specified on the EVNTID parameter in SMEVNTxx.

Default: None

Usage Notes

SUPJSCAT provides jobcat/stepcat suppression for all selected jobs.

Except for EVENTID, there are no rule list parameters for SUPJSCAT. Resource selection is accomplished with the filter list.

Due to the point at which SUPJSCAT is invoked, only the following filter list parameters are available:

DD	DSN	DSNn
HLQ	JOB	LLQ
NQUAL	STEP	UNIT
FUNCTION	RACF	USER
USRCn	USRNn	

Messages

The message issued by SUPJSCAT when a JOBCAT or STEPCAT is suppressed is:

SVM3420I *job,step,dd,dsn* STEPCAT/JOBCAT SUPPRESSED

Note that the SUPJSCAT function must be specified in the SMFUNCxx member to be available and must be set active to provide services.

Example

This example shows you how to suppress jobcat/stepcat cards immediately on test jobs and in simulate mode for production jobs. Started task jobs are exempt, as are jobs submitted by system programmers.

SMFUNC51 member

```
SET  NAME=SUPJSCAT  ACTIVE=YES
      MSG=I  SMF=I
      FLST=49
      DESC='OPEN EMPTY DATA SETS'
```

The function is defined and activated. Note that there is no RLST parameter.

SMFLST49 member

```
SET  MODE=INACT
      INC XMODE=STC
      INC USER=SYSP*
```

Exclude started tasks and system-programmers' jobs from SUPJSCAT processing (by including these resources for inactive mode—this could also have been accomplished by excluding these resources from all SET statements below).

```
SET  MODE=SIM
      INC JOBCLASS=P
```

Simulate job/stepcat suppression for production jobs.

```
SET  MODE=ACT
      INC JOB=/'
```

Suppress job/stepcats for all other jobs.

SUPVOLRF—Suppress Volume References

The SUPVOLRF function suppresses all DASD volume references.

Volume references (VOL=REF= on the DD statement in the JCL) prevent extension of a data set onto additional volumes by the MAINVIEW SRM function SPACVOLA. This restriction can result in unnecessary abends for out-of-space situations.

For each job step, SUPVOLRF examines all DASD data set allocations and removes the volume references for data sets selected by the filter list. It also removes any UNIT=AFF specifications.

Rule List Parameters

SET Statement (in member SMRLSTxx)

Parameter	Description
EVENTID=xxxxx	Identifies a user event defined in an SMEVNTxx member

Rule List Parameter Explanations

EVENTID=

Purpose: Specifies the identifier assigned to a user event in SMEVNTxx. This parameter will cause an event to be generated from this function.

Note: If EVENTID= is used on an FLST SET statement with MODE=INACT, the event will still be issued.

Syntax: EVENTID=xxxxx

where xxxxx is the 5-character string specified on the EVNTID parameter in SMEVNTxx.

Default: None

Usage Notes

SUPVOLRF operates on any DASD data sets, new or existing. SUPVOLRF will not operate on TAPE data sets.

Prior to DFSMS 1.3, allocations that included the VOL=REF parameter were not processed by ACS routines. For new allocations, if the referenced data set was SMS-managed, the storage class of the referenced data set was assigned to a new data set and the new data set became SMS-managed. If the referenced data set was not SMS-managed, the new data set also became non-managed. The SUPVOLRF function can be used to ensure all new allocations are processed through ACS, and the possibility of creating a non-managed data set by simply referring to a non-managed data set is avoided.

Effective with DFSMS 1.3, allocations with VOL=REF coded *are* processed by ACS routines. The volume serial of the referenced data set is not provided in the ACS read-only variables &ANYVOL or &ALLVOLS; rather, pseudo volume serial numbers are used to indicate the type of data set referenced. The VOL filter list parameter will contain one of the following values in the SMSACSDC, SMSACSMC, SMSACSSC, SMSACSSG, and SMSACSTE functions if the SUPVOLRF function *does not* suppress the VOL-REF:

REF=NS the referenced data set is not SMS-managed
REF=ST the referenced data set is an SMS-managed tape file
REF=SD the referenced data set an SMS-managed DASD or VIO file

If SUPVOLRF does suppress the VOL=REF, none of the above pseudo volume serial numbers will be found in the ACS read-only variables &ANYVOL or &ALLVOLS nor in the VOL filter list parameter.

If VOL=REF=dsname refers to an uncataloged data set, the job will normally fail with a JCL error. If the VOL=REF is suppressed by the SUPVOLRF function, the JCL error due to the reference to an uncataloged data set will be avoided. However, if the allocation is for an existing data set, a JCL error may still occur; for example, if the data set is not cataloged, no volume serial information is available to complete the allocation. If the allocation is for a new data set, removing the VOL=REF will make the new allocation non-specific.

The volume serial of the referenced data set is available in the REFVOL filter list parameter during SMSACSDC, SMSACSMC, SMSACSSC, SMSACSSG, and SMSACSTE functions. The REFVOL is also available during the DASDPOOL function. In the case of a reference to an uncataloged data set as discussed above, the REFVOL will contain the value NULVRF.

Because the function has no RLST operands to set, all resource selection must be done through the FLST.

Due to the point at which SUPVOLRF is invoked, only the following filter list parameters are available:

DD	DSN	DSNn
FUNCTION	HLQ	JOB
LLQ	NQUAL	STEP
UNIT	USRCn	USRNn

Note that the SUPVOLRF function must be specified in the SMFUNCxx member to be available and must be set active to provide services.

SUPVOLRF generates message SVM0421I when it removes a volume reference:

```
SVM0421I  job,step,dd,dsn VOL=REF SUPPRESSED
```

Example

This example shows you how to suppress volume references for all Accounts Payable (production) data sets. (In this example, Accounts Payable data sets are identified by a high-level qualifier of DIVA and a third qualifier that starts with TR.)

SMFUNCxx member

```
SET  NAME=SUPVOLRF  ACTIVE=YES  The function is defined and activated.
      MSG=I  SMF=N      There is no RLST specification. Do not
      FLST=34           produce any SMF records.
      DESC=' SUPPRESS VOL=REF'
```

SMFLST34 member

```
SET  MODE=ACT      This filter list specification is active.
      INC HLQ=DIVA DSN3=TR*  Selected resources are all data sets
                              with a high-level qualifier of DIVA and a
                              3rd qualifier beginning with TR.
```

Example

This example shows you how to suppress volume references for a group of frequently selected data set resources identified by a MAINVIEW SRM variable.

SMMSYS10 member

```
SET  FUNC=11  VARS=55  POOL=11      The MAINVIEW SRM system is
      CMDPREF=#  SMFID=251          defined.
      PASSWORD=123456789012
```

SMFUNC11 member

```
SET  NAME=SUPVOLRF  ACTIVE=YES      The function is defined and activated.
      MSG=I  SMF=I      There is no RLST specification.
      FLST=02
      DESC='SUPPRESS VOL=REF'
```

SMFLST02 member

```
SET  MODE=ACT                      This filter list specification is active. It
      INC DSN=&DSNEXT              selects data sets specified in the
                                   variable DSNEXT.
```

SMVARS55 member

```
SET  VARIABLE=DSNEXT               The variable DSNEXT is defined to
      VALUE=AP*.*.MSTR              include the list of data set name
      VALUE=GL035B.TRAN/            masks.
      VALUE=AR*.*.DLYUPD
      VALUE=FLS*/
      VALUE=TPROV?.*.WKLY./
```

TAPECOMP—Control Tape Data Set Compaction (IDRC)

The TAPECOMP function specifies the use of data compaction for a data set on a tape cartridge. Improved Data Recording Capability (IDRC) must be available on the magnetic tape subsystem.

Data compaction on appropriately equipped tape devices can increase the effective capacity by up to 5 times and boost the effective data rate by as much as 70%. This provides the capability of greatly increased overall storage capacity for tape data sets, if it is used consistently. OS/390 requires specification of TRTCH=COMP in the DCB parameters on the DD statement in JCL or the DEVSUPxx member of SYS1.PARMLIB to activate data compaction.

TAPECOMP allows the data center to activate data compaction without changes to JCL. Specification of tape compaction usage can be based on data set names, on device types, on assigned pools, or on any property allowed as a MAINVIEW SRM filter list parameter.

Rule List Parameters

SET Statement (in member SMRLSTxx)

Parameter	Description
COMP=YES/NO	Specifies the use of data compaction
EVENTID=xxxxx	Identifies a user event defined in an SMEVNTxx member

Rule List Parameter Explanations

COMP=

Purpose: Specifies whether data compaction should be used for a data set allocation on a tape cartridge.

Syntax: COMP=YES/NO

Default: None

EVENTID=

Purpose: Specifies the identifier assigned to a user event in SMEVNTxx. This parameter will cause an event to be generated from this function.

Note: If EVENTID= is used on an FLST SET statement with MODE=INACT, the event will still be issued.

Syntax: EVENTID=xxxxx

where xxxxx is the 5-character string specified on the EVNTID parameter in SMEVNTxx.

Default: None

Parameters Not Supported

The following filter list parameters are not supported for TAPECOMP:

ALCTYPE	BUFSP	CATALOG
CONTIG	CURSPACE	DSORG
ENVIR	ERASE	EXTENT
GDGVER	IMBED	MAXSIZE
OWNER	PRISPACE	RECORG
REPL	REUSE	ROUND
SECSPACE	SIZE	VALUE
VSAMCOMP	VSAMDEF	

Usage Notes

TAPECOMP provides services only for tape data sets.

TAPECOMP is not able to verify that the tape device truly supports IDRC.

Note that the TAPECOMP function must be specified in the SMFUNCxx member to be available and must be set active to provide services.

Messages

TAPECOMP generates message SVM3452I when it processes a data set:

SVM3452I *job,step,dd,dsn* TAPE COMPACTION SET TO *on/off*

Example

This example shows you how to activate data compaction for all tape data sets on 3490 devices with IDRC (defined in a pool named 3490C).

SMMSYS10 member

```
SET  FUNC=21 POOL=31
      CMDPREF=# SMFID=251
      PASSWORD=123456789012
```

The MAINVIEW SRM system is defined.

SMFUNC21 member

```
SET  NAME=TAPECOMP
      ACTIVE=YES
      MSG=E SMF=N
      FLST=02 RLST=02
      DESC='COMPACTION FOR
            TAPES'
```

The function is defined and activated. SMF records are not written. Only error messages are written.

SMPOOL31 member

```
SET  POOLNAME=3490C
      TYPE=3490
      INC ADR=180
      ADR=181
      ADR=182
      ADR=183
```

The pool 3490C is defined to include a list of 3490 devices.

SMFLST02 member

```
SET  MODE=ACT
      INC POOL=3490C
```

This filter list specification is active. It selects all data sets going to pool 3490C.

SMRLST02 member

```
SET  COMP=YES
      EXC DISP2=DELETE
```

This rule list specifies data compaction for all data sets whose normal termination disposition is not DELETE.

TAPedefr—Defer Tape Mounts

The TAPedefr function causes OS/390 to defer the tape mount request until the data set is opened.

Data sets allocated to tape devices normally cause a tape mount as soon as the data set is allocated (during job initialization). In many cases, this may be undesirable, since the data set may never be opened. Nevertheless, once the tape mount is issued, the job waits for the device to be readied.

TAPedefr allows the tape mount to be deferred until data set is opened for selected resources.

Rule List Parameters

SET Statement (in member SMRLSTxx)

Parameter	Description
EVENTID=xxxxx	Identifies a user event defined in an SMEVNTxx member

Rule List Parameter Explanations

EVENTID=

Purpose: Specifies the identifier assigned to a user event in SMEVNTxx. This parameter will cause an event to be generated from this function.

Note: If EVENTID= is used on an FLST SET statement with MODE=INACT, the event will still be issued.

Syntax: EVENTID=xxxxx

where xxxxx is the 5-character string specified on the EVENTID parameter in SMEVNTxx.

Default: None

Parameters Not Supported

The following filter list parameters are not supported for TAPedefr:

ALCTYPE	BUFSP	CATALOG
CONTIG	CURSPACE	DSORG
ENVIR	ERASE	EXTENT

GDGVER	IMBED	MAXSIZE
OWNER	PRISPACE	RECORD
REPL	REUSE	ROUND
SECSPACE	SIZE	VALUE
VSAMCOMP	VSAMDEF	

Usage Notes

TAPedefr defers the tape mount request by effectively adding the DEFER subparameter to the UNIT JCL parameter. Although the tape mount is deferred until the data set is opened, the tape device is still allocated at job initialization.

Since there are no rule list parameters for TAPedefr, there is no rule list. Resource selection is accomplished solely with the filter list.

The message issued by TAPedefr when a data set is marked for a deferred mount is:

SVM3454I *job,step,dd,dsn* DEFER ASSIGNED

Note that the TAPedefr function must be specified in the SMFUNCxx member to be available and must be set active to provide services.

Example

This example shows you how to defer tape mounts on all jobs except certain high speed tape related production.

SMFUNC04 member

```
SET  NAME=TAPedefr
      ACTIVE=YES
      MSG=I SMF=I
      FLST=14
      DESC='DEFER TAPE MOUNTS'
```

The function is defined and activated.
Note that there is no RLST parameter.

SMFLST14 member

```
SET  MODE=ACT
      EXC PGM=TSET/
      EXC PGM=TDUP/
      INC JOB=/'
```

Exclude two sets of programs by partial name. For all other jobs, defer tape mounts until the data set is opened.

TAPEPOOL—Assign Tape Data Sets to Pools

The TAPEPOOL function assigns tape data sets to a pool of devices (including DASD devices), overriding any UNIT assignment in the JCL.

The ability to define pools of tape devices, and direct data sets to those pools, allows the storage administrator to more effectively manage the data center's tape resources by assigning usage based on flexible criteria. Long-duration jobs can be assigned to one pool, jobs requiring higher speed can be assigned to a pool containing only 3490 devices, jobs producing very large data sets can be assigned to a pool with IDRC support for compaction, etc.

TAPEPOOL allows data sets to be assigned to pools based on any of the selection criteria supported by MAINVIEW SRM, without reliance on JCL parameters. TAPEPOOL also allows data sets to be redirected from a tape device to a DASD device.

System Parameters

The system parameters that affect TAPEPOOL are described in Table 3-13.

Table 3-13 **System parameters—TAPEPOOL**

Parameter	Purpose
DMYUNIT	Defines the conversion of a nonexistent UNIT parameter to a valid UNIT parameter only if JCLEXT=YES.
REJECT	Controls termination of processing at the first or last data set rejected by REJECT=YES in DASDPOOL or DSNCHECK. If termination is to take place on the first rejected data set, code REJECT=FIRST. If all data sets are to be processed before control is returned to OS/390 allocation, code REJECT=LAST.
TAPEGENR	Specifies tape device generic names that some EasyPOOL functions will intercept. If you want to intercept all tape requests, specify ALLTAPE as the first generic name. (However, you cannot use ALLTAPE when JCLEXT=NO.) TAPEGENR affects functions DSNCHECK, FORCECAT, SETEXPDT, and TAPEPOOL.

Rule List Parameters

SET Statement (in member SMRLSTxx)

Parameter	Description
EVENTID=xxxxx	Identifies a user event defined in an SMEVNTxx member
NVOL=nn	Number of volumes and units that can be allocated to a data set.
POOL=(xxxx,xxxx, . . .)	Specifies one or more pools eligible for allocation
REJECT=YES/NO	Specifies whether the request should be rejected

Parameter	Description
UNIT=xxxxxxx	Unit generic name (1–8 characters)
USEVOL=STOR/PRIV/ALL	Directs volume allocation to a specific type (STOR, PRIV, ALL)

Rule List Parameter Explanations

EVENTID=

Purpose: Specifies the identifier assigned to a user event in SMEVNTxx. This parameter will cause an event to be generated from this function.

Note: If EVENTID= is used on an FLST SET statement with MODE=INACT, the event will still be issued.

Syntax: EVENTID=xxxxx

where xxxxx is the 5-character string specified on the EVNTID parameter in SMEVNTxx.

Default: None

NVOL=

Purpose: Specifies the number of volumes and units that can be allocated to a data set.

Syntax: NVOL=nn

where nn is a number in the range of 1–59 for non-VSAM data sets and 1–20 for VSAM data components.

Default: None

POOL=

Purpose: Specifies one or more pools (defined in SMPOOLxx) to which the data set can be assigned. If the poolname is not defined, the refresh or start-up will fail. This is a required parameter if REJECT=YES is not specified.

Syntax: POOL=(xxxxxx,xxxxxx, . . .)

where xxxxxx specifies pool names from SMPOOLxx. If a single pool is specified, the parentheses are not required.

Default: None

REJECT=

Purpose: Specifies that the request should be rejected; the allocation fails.

Syntax: REJECT=*YES/NO*

Default: REJECT=NO

UNIT=

Purpose: Specifies a unit name (esoteric, generic, or pool name) to be used by the data set during allocation.

Syntax: UNIT=*xxxxxxxx*

where *xx* is a valid unit name (esoteric, generic, or pool name)
1–8 characters long.

Default: None

USEVOL=

Purpose: Specifies the type of volume (storage, private, or either) that activates non-specific allocation requests.

Syntax: USEVOL=*STOR/PRIV/ALL*

STOR storage mounted volume
PRIV privately mounted volume
ALL volume of any mount type

Default: USEVOL=ALL

Parameters Not Supported

The following filter list parameters are not supported for TAPEPOOL:

BUFSP	CATALOG	ENVIR
ERASE	EXTENT	GDGVER
IMBED	OWNER	RECORG
REPL	REUSE	VALUE
VSAMCOMP	VSAMDEF	

Usage Notes

TAPEPOOL applies to data sets assigned to a tape device. The rule list can assign the data set to one or more eligible pools, or reject the allocation request entirely. The pools are defined in the SMPOOL_{xx} PARMLIB member.

TAPEPOOL does not change the device assignment if UNIT=AFF has been specified, or if the data set is a passed data set.

Unless REJECT=YES is specified, the POOL parameter must be specified; TAPEPOOL takes no action without a rule specifying a device pool to which selected resources may be assigned.

TAPEPOOL allows tape data set allocations to be directed to a pool of tape devices defined in SMPOOL_{xx} using the ADR parameter. Allocations to a tape generic or esoteric name, but not a specific unit address, will be allowed to one of the tape units in the pool.

TAPEPOOL allows a data set to be redirected from tape to DASD by specifying a DASD pool in the POOL rule list parameter. If this feature is used, SPACSQTY should also be used to provide space parameters for the DASD allocation; without some primary allocation amount, the DASD allocation for the redirected data set will fail. Note that if the JCL requests a specific tape volume (VOL=SER=xxx), the allocation is not interrupted and, therefore, is not redirected to a DASD pool.

TAPEPOOL ignores DFSMS-managed data sets.

TAPEPOOL does not support multiple pools. They can be specified, but the selected address can be from any one of the pools.

Messages

TAPEPOOL issues these messages:

When a pool is assigned:

SVM3338I *job,step,dd,dsn* ELIGIBLE POOL NAMES SET TO: pools

When an assignment is rejected:

SVM3311I *job,step,dd,dsn* REJECT = reject SET

Note that the TAPEPOOL function must be specified in the SMFUNC_{xx} member to be available and must be set active to provide services.

Example

This example shows you how to:

- Define a tape pool for production jobs only.
- Direct data sets for jobs using UNIT=SYSTAPE to pools containing 3480 and 3490 devices.
- Reject tape data set allocation from users with a TD prefix on their ID's.
- The default tape pool should be 3420 devices.

SMPOOL05 member

SET POOLNAME=T3420 TYPE=3420 INC ADR=(0340,0341)	Define a tape pool of 3420 devices.
SET POOLNAME=T3480 TYPE=3480 INC ADR=(0380,0381,0382)	Define a tape pool of 3480 devices.
SET POOLNAME=T3490A TYPE=3490 INC ADR=(0390,0391)	Define a tape pool of 3490 devices.
SET POOLNAME=T3490B TYPE=3490 INC ADR=(039A,039B,039C,039D)	Define a second tape pool of 3490 devices.
SET POOLNAME=REDIR TYPE=DASD INC VOL=(BNW002, BNW010)	Define a pool of DASD devices for redirection of tape data sets.

SMFUNC02 member

SET NAME=TAPEPOOL ACTIVE=YES MSG=I SMF=N FLST=15 RLST=22 DESC='ASSIGN TAPE POOLS'	The function is defined and activated. The default for SMF records is none.
--	---

SMFLST15 member

SET MODE=ACT INC DSN=/	All data sets are selected for TAPEPOOL processing. (TAPEPOOL ignores allocations other than tape data sets.)
---------------------------	---

SMRLST22 member

SET REJECT=YES INC USER=TD/	Users with IDs beginning with TD are not allowed to run tape jobs.
SET POOL=T3490B INC JOBCLASS=P	Production jobs are assigned to pool T3490B.
SET POOL=(T3480,T3490A) INC UNIT=SYSTAPE	Any data sets originally assigned to the generic unit SYSTAPE will be assigned to a device in pool T3480 or T3490A.
SET POOL=T3420 INC DSN=/	All other tape data sets are sent to the T3420 and T3480 pools.


```
SET POOL=REDIR
INC DSN=HR*.*.CLASS*
INC DSN=HR*.*.PAY*
```

Certain files from the HR system are redirected to a DASD pool. Note that SPACSQTY must also be specified to provide space allocation information.

SMRLST25 member (for SPACSQTY)

```
SET PQTY=100 SQTY=10
INC POOL=REDIR
```

Use SPACSQTY to set a primary and a secondary allocation for all data sets redirected from tape to DASD.

VIOALLOC—Control VIO Allocation

The VIOALLOC function changes the allocation of temporary data sets to virtual storage (VIO).

The OS/390 VIO facility provides a significant advantage over normal DASD I/O. VIO resides in the paging space and I/O operations are performed at the speed of page access.

VIOALLOC changes the allocation of selected temporary data sets to (or from) VIO, with no JCL modifications.

Rule List Parameters

SET Statement (in member SMRLSTxx)

Parameter	Description
EVENTID=xxxxx	Identifies a user event defined in an SMEVNTxx member
VIO=YES/NO	Specifies whether the data set is allocated in VIO space

Rule List Parameter Explanations

EVENTID=

Purpose: Specifies the identifier assigned to a user event in SMEVNTxx. This parameter will cause an event to be generated from this function.

Note: If EVENTID= is used on an FLST SET statement with MODE=INACT, the event will still be issued.

Syntax: EVENTID=xxxxx

where xxxxx is the 5-character string specified on the EVNTID parameter in SMEVNTxx.

Default: None

VIO=

Purpose: Specifies whether the new data set will be allocated in VIO. This is a required parameter; VIO=YES or VIO=NO must be specified for any processing to occur.

Syntax: VIO=*YES/NO*

Default: None

Parameters Not Supported

The following filter list parameters are not supported for VIOALLOC:

- GDGVER
- VALUE

Usage Notes

VIOALLOC applies to new temporary data sets. VIOALLOC changes the allocation of selected data sets to VIO from DASD if the VIO parameter is YES, or to DASD from VIO if the VIO parameter is NO.

VIOALLOC only operates on temporary data sets (DSN=&&xxxx).

MAINVIEW SRM makes the data set eligible for VIO allocation, but OS/390 actually processes the request. Thus, OS/390 considerations apply to data sets switched to VIO. VIO allocation is discussed in “Data Set Resources—Allocation” in the *IBM OS/390 JCL User Guide*.

Messages

VIOALLOC issues two messages.

When a data set is directed to VIO:

SVM3439I *job,step,dd,dsn* MADE VIO ELIGIBLE

When a data set is prohibited from VIO:

SVM3440I *job,step,dd,dsn* VIO ELIGIBILITY REMOVED

If a data set is already allocated VIO and VIO=YES is specified, no action is taken and a message is not issued. Similarly, if a data set is not allocated VIO and VIO=NO is specified, no action is taken and a message is not issued.

Note that the VIOALLOC function must be specified in the SMFUNCxx member to be available and must be set active to provide service.

Example

This example shows you how to:

- Specify VIO usage for all data sets with a maximum allocation of less than five megabytes.
- Ensure sort work data sets are not allocated to VIO.

SMFUNC48 member

SET	NAME=VIOALLOC	ACTIVE=YES	The function is defined and
	MSG=I	SMF=I	activated.
	FLST=21	RLST=54	
	DESC='CONTROL VIO ALLOCATION'		

SMFLST21 member

SET	MODE=ACT	Select all data sets except those
	EXC DDNAME=SORT/	with sort-related ddnames.
	INC DSN=/	

SMRLST54 member

SET	VIO=YES	Allow VIO for data sets whose
	INC MAXSIZE<5MB	maximum size is less than five
		megabytes.

VSAMCNTL—Enforce VSAM Allocation Standards

The VSAMCNTL function allows specification or suppression of various VSAM cluster definitions.

Some parameters that can be specified for VSAM file definitions can have significant affects on processing throughput. Parameters, such as ERASE and IMBED, may be undesirable except under certain circumstances. Other parameters, such as MASTERPW, can affect availability of the data and may be undesirable at your data center.

VSAMCNTL gives the storage administrator control over many of the sensitive VSAM definition parameters. VSAM definition parameters can be overridden for individual data sets or large groups of data sets, based on data set name, job class, user ID, and so on, giving the storage administrator complete control over critical VSAM specifications.

Rule List Parameters

SET Statement (in member SMRLSTxx)

Parameter	Description
BUFSP=nnnnnn	Buffer space for VSAM data sets (0–16776704)
CATALOG=YES/NO	Allows or removes IDCAMS CATALOG parameter
CISIZE=(nnnnnn,nnnnnn)	VSAM data and/or index control interval size (0–999999)
ERASE=YES/NO	Allows or removes IDCAMS ERASE parameter
EVENTID=xxxxx	Identifies a user event defined in an SMEVNTxx member
IMBED=YES/NO	Allows or removes IDCAMS IMBED parameter
OWNER=xxxxxxxxxx	Assigns an owner ID to a VSAM cluster (1–40 characters)
PWDDEL=YES/NO	Delete passwords specified in control statements
REPL=YES/NO	Allows or removes IDCAMS REPLICATE parameter
REUSE=YES/NO	Allows or removes IDCAMS REUSE parameter

Rule List Parameter Explanations**BUFSP=**

Purpose: Contains the buffer space for the cluster or the data component of VSAM data sets. Any existing buffer space specification is overridden.

Note: A BUFSP value that is too small will be ignored by VSAM. Generally, a value less than {2 x data CFSIZE} for non-indexed files or a value less than {2x data CFSIZE+1x index CFSIZE} for indexed files will be too small.

Syntax: BUFSP=nnnnnn

where nnnnnn is a number in the range 0–16776704.

Default: None

CATALOG=

Purpose: Specifies the removal of the CATALOG parameter during cluster definition.

Syntax: CATALOG=YES/NO

If YES is specified the CATALOG parameter is forced in the DEFINE CLUSTER definition; if NO is specified, the CATALOG parameter is removed from the DEFINE CLUSTER definition.

Default: None

CFSIZE=

Purpose: Specifies the value of the data or cluster and/or index control interval size in the corresponding component of the IDCAMS control card(s). Any existing specification of control interval size is overridden.

Syntax: CFSIZE=nnnnnn or CFSIZE=(nnnnnn,nnnnnn)

where the first nnnnnn is the data or cluster control interval size and the second nnnnnn is the index control interval size (0–999999).

Default: None

EVENTID=

Purpose: Specifies the identifier assigned to a user event in SMEVNTxx. This parameter will cause an event to be generated from this function.

Note: If EVENTID= is used on an FLST SET statement with MODE=INACT, the event will still be issued.

Syntax: EVENTID=xxxxx

where xxxxx is the 5-character string specified on the EVNTID parameter in SMEVNTxx.

Default: None

ERASE=

Purpose: Specifies the removal of the ERASE parameter during cluster definition.

Syntax: ERASE=YES/NO

If YES is specified the ERASE parameter is forced in the DEFINE CLUSTER definition; if NO is specified, the ERASE parameter is removed from the DEFINE CLUSTER definition.

Default: None

IMBED=

Purpose: Specifies the removal of the IMBED parameter during cluster definition.

Syntax: IMBED=YES/NO

If YES is specified the IMBED parameter is forced in the DEFINE CLUSTER definition; if NO is specified, the IMBED parameter is removed from the DEFINE CLUSTER definition.

Default: None

OWNER=

Purpose: Specifies an owner to be assigned during cluster definition.

Syntax: OWNER=xxxxxxxxx...

where xxxxxxxxxxx is a string of up to 8 characters or a
MAINVIEW SRM keyword of up to 11 characters.

Default: None

PWDDEL=

Purpose: Specifies that any passwords specified for a VSAM definition be deleted.

Syntax: PWDDEL=YES/NO

The PWDDEL=YES is specified if all password specifications (CONTROLPW, MASTERPW, READPW, UPDATEPW) are removed from the cluster definition; if PWDDEL=NO is specified, any existing password specification is retained.

Default: PWDDEL=NO

REPL=

Purpose: Specifies the removal of the REPLICATE parameter during cluster definition.

Syntax: REPL=YES/NO

If YES is specified the REPLicate parameter is forced in the DEFINE CLUSTER definition; if NO is specified, the REPLicate parameter is removed from the DEFINE CLUSTER definition.

Default: None

REUSE=

Purpose: Specifies the removal of the REUSE parameter during cluster definition.

Syntax: REUSE=YES/NO

If YES is specified, the REUSE parameter is forced in the DEFINE CLUSTER definition; if NO is specified, the REUSE parameter is removed from the DEFINE CLUSTER definition.

Default: None

Parameters Not Supported

The following filter list parameters are not supported for VSAMCNTL:

- GDGVER
- VALUE

Usage Notes

The VSAM definition parameters ERASE, IMBED, and REPLICATE can be removed from a VSAM definition by specifying NO in the rule list. Alternatively, these VSAM parameters can be added to the VSAM file definition by specifying YES.

Note: The used space of the index component will increase in size when IMBED is removed.

The CATALOG=NO rule list parameter removes any CATALOG specification from the VSAM definition; if YES is specified, any existing catalog specification is retained.

The PWDDEL=YES rule list parameter directs that all VSAM definition passwords (CONTROLPW, MASTERPW, READPW, UPDATEPW) are removed from the VSAM definition. If NO is specified, any existing password specifications are retained.

The BUFSP and CISIZE rule list parameters allow the corresponding VSAM parameters to be overridden with the value specified on the MAINVIEW SRM rule list.

The OWNER parameter can be set to a 1–8 byte string that will be assigned during cluster definition. It can also be set to a 1–11 byte MAINVIEW SRM keyword whose value will be used as the owner identification.

Example

In this example

```
SET OWNER=USER
```

will direct MAINVIEW SRM to obtain the RACF user ID and use it as the owner ID.

SET OWNER=JOB

will direct MAINVIEW SRM to use the jobname as the owner ID.

If the length of value to be used for the owner identification exceeds 8 bytes, the remaining data is truncated.

The messages issued by VSAMCNTL are:

SVM3345I *job,step,dd,dsn* REUSE SET TO *yes/no*
SVM3393I *job,step,dd,dsn* CISIZE CHANGED FROM *old* TO *new*
SVM3395I *job,step,dd,dsn* ERASE SET TO *yes/no*
SVM3396I *job,step,dd,dsn* REPLICATE SET TO *yes/no*
SVM3397I *job,step,dd,dsn* IMBED SET TO *yes/no*
SVM3394I *job,step,dd,dsn* BUFSP SET TO *nnnnn*
SVM3398I *job,step,dd,dsn* CATALOG SET TO *yes/no*
SVM3391I *job,step,dd,dsn* PWDDEL SET TO *yes/no*
SVM3392I *job,step,dd,dsn* OWNER SET TO *xxxxx...*

Example

This example shows you how to:

- Remove ERASE, REPLICATE, IMBED, and REUSE from all new VSAM cluster definitions and set the owner identification equal to the RACF user ID.
- Allow only systems programmers the ability to specify CATALOG or passwords.

SMFUNCA1 member

```
SET  NAME=VSAMCNTL ACTIVE=YES
      MSG=I SMF=I FLST=VC RLST=VC
      DESC='VSAM CONTROL'
```

The function is defined and active.

SMFLSTVC member

```
SET  MODE=ACT
      INC JOB=/'
```

Execute the function for all jobs.

SMRLSTVC member

```
SET  ERASE=NO REPL=NO IMBED=NO REUSE=NO X
      EXC OWNER=USER CATALOG=NO PWDDEL=YES
      INC USER=SYSP/
      JOB=/'
```

Disallow the use of ERASE, REPLICATE, IMBED, REUSE, CATALOG, and passwords for non-system programmers. Set the owner ID to the RACF user ID.

```
SET  ERASE=NO REPL=NO IMBED=NO REUSE=NO X
      INC OWNER=USER
      JOB=/'
```

Disallow the use of ERASE, REPLICATE, IMBED and REUSE for system programmers. Set the owner ID to the RACF user ID.

Chapter 4 How to Define Critical Data Set Lists

EasyPOOL allows you to identify a list of critical data sets in your environment that should never be placed on the same volume as any other data set in the list. It is important to consider, however, that CRITDSN is resource intensive and should be used only for a *small* list of critical data sets. It should not be used without considering the system impact.

To define critical lists, you must

- define a critical data set list (see “How to Use the CritList Option” on page 4-3 for instructions) and add it to the RLST
- define and activate the function(s) (DASDPOOL or SMSSELECT) with the critical data set list active
- activate the system member that includes the critical data set list information

You can change SMCRTxx parameters

- by editing the member directly
- by using the MAINVIEW SRM function SET statements
- through the **CritList** option on the Parmlib Members menu (see “How to Use the CritList Option” on page 4-3)

Example

This example shows you how critical data set list information is defined in each of the members. If you need further instructions, refer to the *MAINVIEW SRM User Guide and Reference*.

SMCRIT01 member		
SET	CRI TDSN=PROD01 INC DSN=PROD. PARTS. DATABASE INC DSN=PROD.ASSEMBLE.DATABASE INC DSN=PROD.MASTER.DATABASE	Critical data set list PROD01 is created in SMCRIT01.
SMRLSTVC member		
SET	MODE=ACT VOLSEL=CRI TDSN CRITBIAS=10 CRITFAIL=NO CRITLIST=PROD01 INC DSN=PROD.1	Critical data set list PROD01 is specified in the SMRLSTVC member along with other critical data set list processing parameters.
SMFLSTVC member		
SET	MODE=ACT INC JOB=PRODJOB1	Critical data set list PROD01 is specified in the SMRLSTVC member along with other critical data set list processing parameters.
SMFUNCAC member		
SET	NAME=DASDPOOL ACTIVE=YES MSG=I SMF=I FLST=VC RLST=VC DESC='DASD ALLOCATION'	The function is defined and active.
SMMSYS75 member		
SET	CRITLIST=01 FUNC=AC	The critical list suffix from SMCRIT01 is defined in SMMSYS75 as CRITLIST=01. the SMFUNCAC member is defined as FUNC=AC.

How to Use the CritList Option

Once you have identified the data sets that are critical at your data center, you may use the Critlist (EasyPOOL) option to define critical data set lists as instructed in the next table.

Step	Action
1.	Select Parmlib Members on the EZSRM Menu. A list of members displays.
2.	Select CritList from the selection list.
3.	Type E in the line command field of the CMD column next to an active SMCRTxx member or place your cursor on the member name, then press Enter .
4.	Edit the member to specify the data center critical data sets.
5.	Press PF3 to save and exit the member.
6.	Type R in the line command field of the CMD column next to the SMCRTxx member you edited, then press Enter . The member is refreshed. Note: By refreshing SMCRTxx, you have made it the active critical list. It will remain active until SVOS is stopped or until you refresh another critical list member.

Chapter 5 How to Define DFSMS Subpools

DFSMS-managed subpooling is a method by which DFSMS-managed Storage Groups, which are defined in ACS code and cannot be changed, can be subdivided into pools to improve data set allocation and performance. This functionality is limited to denying volumes outside of an DFSMS-defined pool. Volume selection cannot be directed outside of the data set's current Storage Group. If multiple pools are specified, the first pool with sufficient space for the data set is assigned. The first pool specified is the primary pool; all others are alternate pools.

To define DFSMS-managed subpools:

- Define the subpools within a Storage Group using the SMSPOLxx definition parameter member (see “How to Use the SMS Pools Option” on page 5-5).
- Activate the SMSSELECT function in EasyPOOL using the SMFUNCxx definition parameter member (refer to the *MAINVIEW SRM User Guide and Reference* for instructions).
- Set system options SMS_ALLOC and SMS_EXTEND to Yes (refer to the *MAINVIEW SRM User Guide and Reference* for instructions).
- Assign SMSPOOLS to data sets with the parameters SMSPOOL and SMSPOOL_EXT (refer to the *MAINVIEW SRM User Guide and Reference* for instructions).

SMSPOL_{xx}

SMSPOL_{xx} is the DFSMS-managed subpool definition parameter member. It is read by MAINVIEW SRM during system start up. The suffix specification in SMMSYS_{xx} indicates which version of SMSPOL_{xx} contains SMS-managed subpool specifications for the particular configuration of MAINVIEW SRM that is being executed.

Parameters

Subpools are defined on SET statements. INC/EXC statements that follow a SET statement identify the devices to be included or excluded from a subpool.

You can change SMSPOL_{xx} parameters

- by editing the member directly
- by using the MAINVIEW SRM system function SET statements
- through the **SMS Pools** option on the EZSRMP Parmlib Members menu (see “How to Use the SMS Pools Option” on page 5-5)

Parameter Quick Reference

SET statement pool parameters are described in the following table. Detailed descriptions of parameters can be found in the *MAINVIEW SRM User Guide and Reference*.

Table 5-1 SET Statement SMS-Managed Subpool Parameter Quick Reference

Parameter	Required	Description
POOLNAME=xxxxxxx	Yes	Name to be assigned to a pool

INC/EXC statement for SMS-managed subpool parameters are described in the following quick reference table.

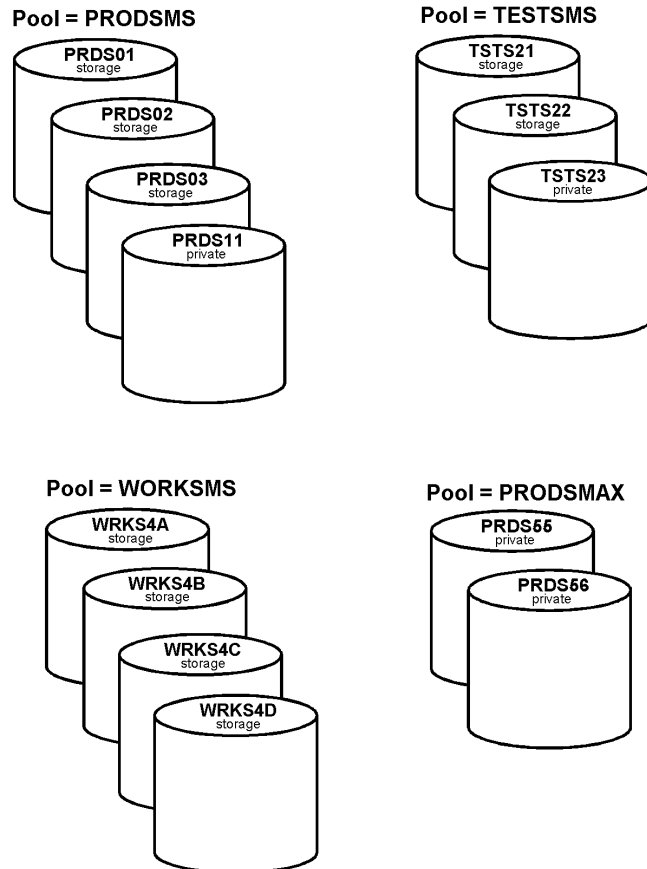
Table 5-2 INC/EXC Statement Pool Parameter Quick Reference

Parameter	Required	Description
VOL=xxxxxx	No	Volume serial number of device in pool
VOL=(xxxxxx,xxxxxx,...)	No	Multiple volume serial numbers (up to 15)

Example

This example shows you how to configure DFSMS-managed subpools using the information given in the following figure.

Figure 5-1 Possible System Storage Configuration for SMS-Managed Subpools



SMSPOL00 member	
SET POOLNAME=PRODSMS INC VOL=PRDS0/ INC VOL=PRDS11	The PRODSMS subpool is composed of all volumes whose names begin with PRDS0 and volume PRDS11.
SET POOLNAME=PRODSMAX INC VOL=(PRDS55,PRDS56)	The PRODSMAX subpool contains volumes PRDS55 and PRDS56.
SET POOLNAME=TESTSMS INC VOL=(TSTS21,TSTS22,TSTS23)	The TESTSMS subpool is composed of volumes TSTS21, TSTS22, and TSTS23.
SET POOLNAME=WORKSMS INC VOL=WRKS4/	The WORKSMS subpool contains all volumes whose names begin with WRKS4.

How to Use the SMS Pools Option

Once you have identified the criteria for SMS-managed subpools at your data center, you may use the SMS Pools option to define pools as instructed in the next table.

Step	Action
1.	Select Parmlib Members on the EZSRM Menu. A list of members displays.
2.	Select SMS Pools from the selection list.
3.	Type E in the line command field of the CMD column next to an active SMSPOLxx member or place your cursor on the member name, then press Enter .
4.	Edit the member to specify the data center pool definitions.
5.	Press PF3 to save and exit the member.
6.	Type R in the line command field of the CMD column next to the SMPOOLxx member you edited, then press Enter . The member is refreshed. Note: By refreshing SMSPOLxx, you have made it the active SMS-managed subpool list. It will remain active until SVOS is stopped or until you refresh another pool member.

Glossary

This glossary defines BMC Software terminology. Other dictionaries and glossaries can be used in conjunction with this glossary.

Since this glossary pertains to BMC Software-related products, some of the terms defined might not appear in this book.

To help you find the information you need, this glossary uses the following cross-references:

Contrast with indicates a term that has a contrary or contradictory meaning.

See indicates an entry that is a synonym or contains expanded information.

See also indicates an entry that contains related information.

action	Defined operation, such as modifying a MAINVIEW window, that is performed in response to a command. <i>See</i> object.
active window	Any MAINVIEW window in which data can be refreshed. <i>See</i> alternate window, current window, window.
administrative view	Display from which a product's management tasks are performed, such as the DSLIST view for managing historical data sets. <i>See</i> view.
ALT WIN field	Input field that allows you to specify the window identifier for an alternate window where the results of a hyperlink are displayed. <i>See</i> alternate window.
Alternate Access	<i>See</i> MAINVIEW Alternate Access.
alternate form	View requested through the FORM command that changes the format of a previously displayed view to show related information. <i>See also</i> form, query.

alternate window	(1) Window that is specifically selected to display the results of a hyperlink. (2) Window whose identifier is defined to the ALT WIN field. <i>Contrast with</i> current window. <i>See</i> active window, window, ALT WIN field.
analyzer	(1) Online display that presents a snapshot of status and activity data and indicates problem areas. (2) Component of CMF MONITOR. <i>See</i> CMF MONITOR Analyzer.
application	(1) Program that performs a specific set of tasks within a MAINVIEW product. (2) In MAINVIEW VistaPoint, combination of workloads to enable display of their transaction performance data in a single view.
application trace	<i>See</i> trace.
ASCH workload	Workload comprising Advanced Program-to-Program Communication (APPC) address spaces.
AutoCustomization	Online facility for customizing the installation of products. AutoCustomization provides an ISPF panel interface that both presents customization steps in sequence and provides current status information about the progress of the installation.
automatic screen update	Usage mode wherein the currently displayed screen is refreshed automatically with new data at an interval you specify. Invoked by the ASU command.
batch workload	Workload consisting of address spaces running batch jobs.
BBI	Basic architecture that distributes work between workstations and multiple OS/390 targets for BMC Software MAINVIEW products.
BBI-SS PAS	<i>See</i> BBI subsystem product address space.
BBI subsystem product address space (BBI-SS PAS)	OS/390 subsystem address space that manages communication between local and remote systems and that contains one or more of the following products: <ul style="list-style-type: none"> • MAINVIEW AutoOPERATOR • MAINVIEW for CICS • MAINVIEW for DB2 • MAINVIEW for DBCTL • MAINVIEW for IMS Online • MAINVIEW for MQSeries (formerly Command MQ for S/390) • MAINVIEW SRM • MAINVIEW VistaPoint (for CICS, DB2, DBCTL, and IMS workloads)
BBPARM	<i>See</i> parameter library.

BBPROC	<i>See</i> procedure library.
BBPROF	<i>See</i> profile library.
BBSAMP	<i>See</i> sample library.
BBV	<i>See</i> MAINVIEW Alternate Access.
BBXS	BMC Software Subsystem Services. Common set of service routines loaded into common storage and used by several BMC Software MAINVIEW products.
border	Visual indication of the boundaries of a window.
bottleneck analysis	Process of determining which resources have insufficient capacity to provide acceptable service levels and that therefore can cause performance problems.
CA-Disk	Data management system by Computer Associates that replaced the DMS product.
CAS	Coordinating address space. One of the address spaces used by the MAINVIEW windows environment architecture. The CAS supplies common services and enables communication between linked systems. Each OS/390 or z/OS image requires a separate CAS. Cross-system communication is established through the CAS using VTAM and XCF communication links.
CFMON	<i>See</i> coupling facility monitoring.
chart	Display format for graphical data. <i>See also</i> graph.
CICSplex	User-defined set of one or more CICS systems that are controlled and managed as a single functional entity.
CMF MONITOR	Comprehensive Management Facility MONITOR. Product that measures and reports on all critical system resources, such as CPU, channel, and device usage; memory, paging, and swapping activity; and workload performance.
CMF MONITOR Analyzer	Batch component of CMF MONITOR that reads the SMF user and 70 series records created by the CMF MONITOR Extractor and/or the RMF Extractor and formats them into printed system performance reports.
CMF MONITOR Extractor	Component of CMF that collects performance statistics for CMF MONITOR Analyzer, CMF MONITOR Online, MAINVIEW for OS/390, and RMF postprocessor. <i>See</i> CMF MONITOR Analyzer, CMF MONITOR Online, MAINVIEW for OS/390.

CMF MONITOR Online

Component of CMF that uses the MAINVIEW window interface to present data on all address spaces, their use of various system resources, and the delays that each address space incurs while waiting for access to these resources. *See* CMF MONITOR, MAINVIEW for OS/390.

CMF Type 79 API

Application programming interface, provided by CMF, that provides access to MAINVIEW SMF-type 79 records.

CMFMON

Component of CMF MONITOR that simplifies online retrieval of information about system hardware and application performance and creates MAINVIEW SMF-type 79 records.

The CMFMON *online facility* can be used to view data in one or more formatted screens.

The CMFMON *write facility* can be used to write collected data as MAINVIEW SMF-type 79 records to an SMF or sequential data set.

CMRDETL

MAINVIEW for CICS data set that stores detail transaction records (type 6E) and abend records (type 6D). Detail records are logged for each successful transaction. Abend records are written when an abend occurs. Both records have the same format when stored on CMRDETL.

CMRSTATS

MAINVIEW for CICS data set that stores both CICS operational statistic records, at five-minute intervals, and other records, at intervals defined by parameters specified during customization (using CMRSOPT).

column

Vertical component of a view or display, typically containing fields of the same type of information, that varies by the objects associated in each row.

collection interval

Length of time data is collected. *See also* delta mode, total mode.

command delimiter

Special character, usually a ; (semicolon), used to stack commands typed concurrently on the COMMAND line for sequential execution.

COMMAND line

Line in the control area of the display screen where primary commands can be typed. *Contrast with* line command column.

Command MQ Automation D/S

Command MQ agents, which provide local proactive monitoring for both MQSeries and MSMQ (Microsoft message queue manager). The Command MQ agents operate at the local node level where they continue to perform functions regardless of the availability of the MQM (message queue manager) network. Functionality includes automatic monitoring and restarts of channels, queue managers, queues and command servers. In cases where automated recovery is not possible, the agents transport critical alert information to a central console.

Command MQ Automation S/390

Command MQ component, which monitors the MQM (message queue manager) networks and intercedes to perform corrective actions when problems arise. Solutions include:

- Dead-Letter Queue management
- System Queue Archival
- Service Interval Performance solutions
- Channel Availability

These solutions help ensure immediate relief to some of the most pressing MQM operations and performance problems.

Command MQ for D/S

Command MQ for D/S utilizes a true client/server architecture and employs resident agents to provide configuration, administration, performance monitoring and operations management for the MQM (message queue manager) network.

Command MQ for S/390

See MAINVIEW for MQSeries.

COMMON STORAGE MONITOR

Component of MAINVIEW for OS/390 that monitors usage and reconfigures OS/390 or z/OS common storage blocks.

composite workload

Workload made up of a WLM workload or other workloads, which are called *constituent workloads*.

constituent workload

Member of a composite workload. Constituent workloads in a composite usually belong to a single workload class, but sometimes are mixed.

contention

Occurs when there are more requests for service than there are servers available.

context

In a Plex Manager view, field that contains the name of a target or group of targets specified with the CONTEXT command. *See* scope, service point, SSI context, target context.

CONTEXT command

Specifies either a MAINVIEW product and a specific target for that product (*see* target context) or a MAINVIEW product and a name representing one or more targets (*see* SSI context) for that product.

control statement	(1) Statement that interrupts a sequence of instructions and transfers control to another part of the program. (2) Statement that names samplers and other parameters that configure the MAINVIEW components to perform specified functions. (3) In CMF MONITOR, statement in a parameter library member used to identify a sampler in the extractor or a report in the analyzer, or to describe either component's processing requirements to the operating system.
coupling facility monitoring (CFMON)	Coupling facility views that monitor the activity of your system's coupling facilities.
current data	Data that reflects the system in its current state. The two types of current data are real-time data and interval data. <i>Contrast with</i> historical data. <i>See also</i> interval data, real-time data.
current window	In the MAINVIEW window environment, window where the main dialog with the application takes place. The current window is used as the default window destination for commands issued on the COMMAND line when no window number is specified. <i>Contrast with</i> alternate window. <i>See</i> active window, window.
DASD	(Direct Access Storage Device) (1) A device with rotating recording surfaces that provides immediate access to stored data. (2) Any device that responds to a DASD program.
DASD ADVISOR	An interactive software tool that diagnoses DASD performance problems and makes recommendations to reduce overall service time. This tool measures and reports on the operational performance of IBM and IBM-compatible devices.
data collector	Program that belongs to a MAINVIEW product and that collects data from various sources and stores the data in records used by views. For example, MAINVIEW for OS/390 data collectors obtain data from OS/390 or z/OS services, OS/390 or z/OS control blocks, CMF MONITOR Extractor control blocks, and other sources. <i>Contrast with</i> extractor.
delta mode	(1) In MAINVIEW for DB2 analyzer displays, difference between the value sampled at the start of the current statistics interval and the value sampled by the current analyzer request. <i>See also</i> statistics interval. (2) In CMFMON, usage mode wherein certain columns of data reflect the difference in values between one sample cycle and the next. Invoked by the DELta ON command. <i>See also</i> collection interval, sample cycle, total mode.
DFSMS	(Data Facility Storage Management System) Data management, backup, and HSM software from IBM for OS/390 or z/OS mainframes.
DMR	<i>See</i> MAINVIEW for DB2.

DMS	(Data Management System) <i>See</i> CA-Disk.
DMS2HSM	<i>See</i> MAINVIEW SRM DMS2HSM.
DSO	(Data Set Optimizer) CMF MONITOR Extractor component that uses CMF MONITOR Extractor data to produce reports specifying the optimal ordering of data sets on moveable head devices.
EasyHSM	<i>See</i> MAINVIEW SRM EasyHSM.
EasyPOOL	<i>See</i> MAINVIEW SRM EasyPOOL.
EasySMS	<i>See</i> MAINVIEW SRM EasySMS.
element	(1) Data component of a data collector record, shown in a view as a field. (2) Internal value of a field in a view, used in product functions.
element help	Online help for a field in a view. The preferred term is <i>field help</i> .
Enterprise Storage Automation	<i>See</i> MAINVIEW SRM Enterprise Storage Automation.
event	A message issued by Enterprise Storage Automation. User-defined storage occurrences generate events in the form of messages. These events provide an early warning system for storage problems and are routed to user-specified destinations for central viewing and management.
Event Collector	Component for MAINVIEW for IMS Online, MAINVIEW for IMS Offline, and MAINVIEW for DBCTL that collects data about events in the IMS environment. This data is required for Workload Monitor and optional for Workload Analyzer (except for the workload trace service). This data also is recorded as transaction records (X'FA') and program records (X'F9') on the IMS system log for later use by the MAINVIEW for IMS Offline components: Performance Reporter and Transaction Accountant.
expand	Predefined link from one display to a related display. <i>See also</i> hyperlink.
extractor	Program that collects data from various sources and keeps the data control blocks to be written as records. Extractors obtain data from services, control blocks, and other sources. <i>Contrast with</i> data collector.
extractor interval	<i>See</i> collection interval.
fast path	Predefined link between one screen and another. To use the fast path, place the cursor on a single value in a field and press Enter . The resulting screen displays more detailed information about the selected value. <i>See also</i> hyperlink.

field	Group of character positions within a screen or report used to type or display specific information.
field help	Online help describing the purpose or contents of a field on a screen. To display field help, place the cursor anywhere in a field and press PF1 (HELP). In some products, field help is accessible from the screen help that is displayed when you press PF1 .
filter	Selection criteria used to limit the number of rows displayed in a view. Data that does not meet the selection criteria is not displayed. A filter is composed of an element, an operator, and an operand (a number or character string). Filters can be implemented in view customization, through the PARM/QPARM commands, or through the Where/QWhere commands. Filters are established against elements of data.
fire	The term used to indicate that an event has triggered an action. In MAINVIEW AutoOPERATOR, when a rule selection criteria matches an incoming event and <i>fires</i> , the user-specified automation actions are performed. This process is also called <i>handling</i> the event.
fixed field	Field that remains stationary at the left margin of a screen that is scrolled either right or left.
FOCAL POINT	MAINVIEW product that displays a summary of key performance indicators across systems, sites, and applications from a single terminal.
form	One of two constituent parts of a view; the other is query. A form defines how the data is presented; a query identifies the data required for the view. <i>See also</i> query, view.
full-screen mode	Display of a MAINVIEW product application or service on the entire screen. There is no window information line. <i>Contrast with</i> windows mode.
global command	Any MAINVIEW window interface command that can affect all windows in the window area of a MAINVIEW display.
graph	Graphical display of data that you select from a MAINVIEW window environment view. <i>See also</i> chart.
hilevel	For MAINVIEW products, high-level data set qualifier required by a site's naming conventions.
historical data	(1) Data that reflects the system as it existed at the end of a past recording interval or the duration of several intervals. (2) Any data stored in the historical database and retrieved using the TIME command. <i>Contrast with</i> current data, interval data and real-time data.

historical database	Collection of performance data written at the end of each installation-defined recording interval and containing up to 100 VSAM clusters. Data is extracted from the historical database with the TIME command. <i>See</i> historical data.
historical data set	In MAINVIEW products that display historical data, VSAM cluster file in which data is recorded at regular intervals.
HSM	(Hierarchical Storage Management) Automatic movement of files from hard disk to slower, less-expensive storage media. The typical hierarchy is from magnetic disk to optical disk to tape.
hyperlink	<p>(1) Preset field in a view or an EXPAND line on a display that permits you to</p> <ul style="list-style-type: none"> • access cursor-sensitive help • issue commands • link to another view or display <p>The transfer can be either within a single product or to a related display/view in a different BMC Software product. Generally, hyperlinked fields are highlighted. (2) Cursor-activated short path from a topic or term in online help to related information. <i>See also</i> fast path.</p>
Image log	<p>Collection of screen-display records. Image logs can be created for both the BBI-SS PAS and the BBI terminal session (TS).</p> <p>The BBI-SS PAS Image log consists of two data sets that are used alternately: as one fills up, the other is used. Logging to the BBI-SS PAS Image log stops when both data sets are filled and the first data set is not processed by the archive program.</p> <p>The TS Image log is a single data set that wraps around when full.</p>
IMSPlex System Manager (IPSM)	MVIMS Online and MVDBC service that provides Single System Image views of resources and bottlenecks for applications across one or more IMS regions and systems.
interval data	<p>Cumulative data collected during a collection interval. Intervals usually last from 15 to 30 minutes depending on how the recording interval is specified during product customization. <i>Contrast with</i> historical data.</p> <p>Note: If change is made to the workloads, a new interval will be started.</p> <p><i>See also</i> current data and real-time data.</p>
InTune	Product for improving application program performance. It monitors the program and provides information used to reduce bottlenecks and delays.

IRUF	IMS Resource Utilization File (IRUF). IRUFs can be either detail (one event, one record) or summarized (more than one event, one record). A detail IRUF is created by processing the IMS system log through a program called IMFLEEDIT. A summarized IRUF is created by processing one or more detail IRUFs, one or more summarized IRUFs, or a combination of both, through a sort program and the TASCOSTR program.
job activity view	Report about address space consumption of resources. <i>See</i> view.
journal	Special-purpose data set that stores the chronological records of operator and system actions.
Journal log	<p>Collection of messages. Journal logs are created for both the BBI-SS PAS and the BBI terminal session (TS).</p> <p>The BBI-SS PAS Journal log consists of two data sets that are used alternately: as one fills up, the other is used. Logging to the BBI-SS PAS Journal log stops when both data sets are filled and the first data set is not being processed by the archive program.</p> <p>The TS Journal log is a single data set that wraps around when full.</p>
line command	Command that you type in the line command column in a view or display. Line commands initiate actions that apply to the data displayed in that particular row.
line command column	Command input column on the left side of a view or display. <i>Contrast with</i> COMMAND line.
Log Edit	In the MAINVIEW for IMS Offline program named IMFLEEDIT, function that extracts transaction (X'FA') and program (X'F9') records from the IMS system log. IMFLEEDIT also extracts certain records that were recorded on the system log by IMS. IMFLEEDIT then formats the records into a file called the IMS Resource Utilization File (IRUF).
MAINVIEW	BMC Software integrated systems management architecture.
MAINVIEW Alarm Manager (MV ALARM)	In conjunction with other MAINVIEW products, notifies you when an exception occurs. MAINVIEW Alarm Manager is capable of monitoring multiple systems simultaneously, which means that MAINVIEW Alarm Manager installed on one system keeps track of your entire sysplex. You can then display a single view that shows exceptions for all MAINVIEW performance monitors within your OS/390 or z/OS enterprise.

MAINVIEW Alternate Access

Enables MAINVIEW products to be used without TSO by providing access through EXCP and VTAM interfaces.

MAINVIEW Application Program Interface (MVAPI)

A CLIST- or REXX-based, callable interface that allows MAINVIEW AutoOPERATOR EXECs to access MAINVIEW monitor product view data.

MAINVIEW AutoOPERATOR

Product that uses tools, techniques, and facilities to automate routine operator tasks and provide online performance monitoring, and that achieves high availability through error minimization, improved productivity, and problem prediction and prevention.

MAINVIEW control area

In the MAINVIEW window environment, first three lines at the top of the view containing the window information line and the COMMAND, SCROLL, CURR WIN, and ALT WIN lines. The control area cannot be customized and is part of the information display. *Contrast with* MAINVIEW display area, MAINVIEW window area.

MAINVIEW Desktop Version of the MAINVIEW window interface designed to run on OS/2 and Windows workstations.

MAINVIEW display area

See MAINVIEW window area.

MAINVIEW Explorer Product that provides access to MAINVIEW products from a Web browser running under Windows. MAINVIEW Explorer replaces MAINVIEW Desktop.

MAINVIEW for CICS Product (formerly MV MANAGER for CICS) that provides real-time application performance analysis and monitoring for CICS system management.

MAINVIEW for DB2 Product (formerly MV MANAGER for DB2) that provides real-time and historical application performance analysis and monitoring for DB2 subsystem management.

MAINVIEW for DBCTL (MVDBC)

Product that provides real-time application performance analysis and monitoring for DBCTL management.

MAINVIEW for IMS (MVIMS) Offline

Product with a Performance Reporter component that organizes data and prints reports used to analyze IMS performance and a Transaction Accountant component that produces cost accounting and user charge-back records and reports.

MAINVIEW for IMS (MVIMS) Online

Product that provides real-time application performance analysis and monitoring for IMS management.

MAINVIEW for IP

Product that monitors OS/390 and z/OS mission-critical application performance as it relates to TCP/IP stack usage. Collected data includes availability, connections, response times, routers, service levels, storage, traffic, Web cache, and so on.

MAINVIEW for Linux–Servers

Product that allows you to monitor the performance of your Linux systems from the MAINVIEW windows interface.

MAINVIEW for MQSeries (formerly known as Command MQ for S/390)

Delivers comprehensive capabilities for configuration, administration, performance monitoring and operations management for an entire MQM (message queue manager) network.

MAINVIEW for OS/390

System management application (formerly known as MAINVIEW for MVS prior to version 2.5). Built upon the MAINVIEW window environment architecture, it uses the window interface to provide access to system performance data and other functions necessary in the overall management of an enterprise.

MAINVIEW for UNIX System Services

System management application that allows you to monitor the performance of the Unix System Services from a MAINVIEW window interface.

MAINVIEW for VTAM

Product that displays application performance data by application, transaction ID, and LU name. This collected data includes connections, response time statistics, application availability, and application throughput.

MAINVIEW for WebSphere Application Server (formerly known as MAINVIEW for WebSphere)

Product that provides extensive monitoring for the IBM WebSphere Application Server for z/OS and OS/390 environment.

MAINVIEW Selection Menu

ISPF selection panel that provides access to all MAINVIEW windows-mode and full-screen mode products.

MAINVIEW SRM

See MAINVIEW Storage Resource Manager (SRM).

MAINVIEW SRM DMS2HSM

Product that facilitates the conversion of CA-Disk, formerly known as DMS, to HSM.

MAINVIEW SRM EasyHSM

Product that provides online monitoring and reporting to help storage managers use DFHSM efficiently.

MAINVIEW SRM EasyPOOL

Product that provides control over data set allocation and enforcement of allocation and naming standards. EasyPOOL functions operate at the operating system level to intercept normal job processing, thus providing services without any JCL changes.

MAINVIEW SRM EasySMS

Product that provides tools that aid in the conversion to DFSMS and provides enhancement to the DFSMS environment after implementation. EasySMS consists of the EasyACS functions, the SMSACSTE function, and the Monitoring and Positioning Facility.

MAINVIEW SRM Enterprise Storage Automation

Product that delivers powerful event generation and storage automation technology across the storage enterprise. Used in conjunction with MAINVIEW AutoOPERATOR, automated solutions to perform pool, volume, application, or data set-level manipulation can be created and used in response to any condition or invoked to perform ad hoc requests.

MAINVIEW SRM SG-Auto

Product that provides early warning notification of storage anomalies and automated responses to those anomalies based on conditions in the storage subsystem.

MAINVIEW SRM SG-Control

Product that provides real-time monitoring, budgeting, and control of DASD space utilization.

MAINVIEW SRM StopX37/II

Product that provides enhancements to OS/390 or z/OS space management, reducing the incidence of space-related processing problems. The StopX37/II functions operate at the system level to intercept abend conditions or standards violations, thus providing services without any JCL changes.

MAINVIEW SRM StorageGUARD

Product that monitors and reports on DASD consumption and provides historical views to help control current and future DASD usage.

MAINVIEW Storage Resource Manager (SRM)

Suite of products that assist in all phases of OS/390 or z/OS storage management. MAINVIEW SRM consists of products that perform automation, reporting, trend analysis, and error correction for storage management.

MAINVIEW SYSPROG Services

See SYSPROG services.

MAINVIEW VistaPoint

Product that provides enterprise-wide views of performance. Application and workload views are available for CICS, DB2, DBCTL, IMS, OS/390, or z/OS. Data is summarized at the level of detail needed; for example, views can be for a single target, an OS/390 or z/OS image, or an entire enterprise.

MAINVIEW window area

Portion of the information display that is not the control area and in which views are displayed and windows opened. It includes all but the first three lines of the information display. *Contrast with* MAINVIEW control area.

monitor

Online service that measures resources or workloads at user-defined intervals and issues warnings when user-defined thresholds are exceeded.

Multi-Level Automation (MLA)

The user-defined, multiple step process in Enterprise Storage Automation that implements solutions in a tiered approach, where solutions are invoked one after another until the condition is resolved.

MVALARM

See MAINVIEW Alarm Manager.

MVAPI

See MAINVIEW Application Program Interface.

MVCICS

See MAINVIEW for CICS.

MVDB2

See MAINVIEW for DB2.

MVDBC

See MAINVIEW for DBCTL.

MVIMS

See MAINVIEW for IMS.

MVIP

See MAINVIEW for IP.

MVLNX

See MAINVIEW for Linux–Servers.

MVMQ

See MAINVIEW for MQSeries.

MVMVS

See MAINVIEW for OS/390.

MVScope

MAINVIEW for OS/390 application that traces both CPU usage down to the CSECT level and I/O usage down to the channel program level.

MVSRM

See MAINVIEW Storage Resource Manager (SRM).

MVSRMHSM

See MAINVIEW SRM EasyHSM.

MVSRMSGC	<i>See</i> MAINVIEW SRM SG-Control.
MVSRMSGD	<i>See</i> MAINVIEW SRM StorageGUARD.
MVSRMSGP	<i>See</i> MAINVIEW SRM StorageGUARD.
MVUSS	<i>See</i> MAINVIEW for UNIX System Services.
MVVP	<i>See</i> MAINVIEW VistaPoint.
MVVTAM	<i>See</i> MAINVIEW for VTAM.
MVWEB	<i>See</i> MAINVIEW for WebSphere Application Server.
nested help	Multiple layers of help pop-up windows. Each successive layer is accessed by clicking a hyperlink from the previous layer.
object	<p>Anything you can manipulate as a single unit. MAINVIEW objects can be any of the following: product, secondary window, view, row, column, or field.</p> <p>You can issue an action against an object by issuing a line command in the line command column to the left of the object. <i>See</i> action.</p>
OMVS workload	Workload consisting of OS/390 OpenEdition address spaces.
online help	Help information that is accessible online.
OS/390 and z/OS Installer	BMC Software common installation system for mainframe products.
OS/390 product address space (PAS)	Address space containing OS/390 or z/OS data collectors, including the CMF MONITOR Extractor. Used by MAINVIEW for OS/390, MAINVIEW for UNIX System Services, and CMF MONITOR products. <i>See</i> PAS.
parameter library	<p>Data set consisting of members that contain parameters for specific MAINVIEW products or a support component. There can be several versions:</p> <ul style="list-style-type: none"> the distributed parameter library, called BBPARM a site-specific parameter library or libraries <p>These can be</p> <ul style="list-style-type: none"> a library created by AutoCustomization, called UBBPARM a library created manually, with a unique name

PAS	Product address space. Used by the MAINVIEW products. Contains data collectors and other product functions. <i>See also</i> OS/390 product address space (PAS) <i>and</i> BBI subsystem product address space (BBI-SS PAS).
performance group workload	Collection of address spaces defined to OS/390 or z/OS. If you are running OS/390 or z/OS with WLM in compatibility mode, MAINVIEW for OS/390 creates a performance group workload instead of a service class.
PERFORMANCE MANAGER	MAINVIEW for CICS online service for monitoring and managing current performance of CICS regions.
Performance Reporter (MVIMS)	MVIMS Offline component that organizes data and prints reports that can be used to analyze IMS performance.
Performance Reporter	Product component that generates offline batch reports. The following products can generate these reports: <ul style="list-style-type: none">• MAINVIEW for DB2• MAINVIEW for CICS
Plex Manager	Product through which cross-system communication, MAINVIEW security, and an SSI context are established and controlled. Plex Manager is shipped with MAINVIEW window environment products as part of the coordinating address space (CAS) and is accessible as a menu option from the MAINVIEW Selection Menu.
pop-up display	Full-screen panel that displays additional information about a selected event in a detail trace.
pop-up window	Window containing help information that, when active, overlays part of the window area. A pop-up window is displayed when you issue the HELP command while working in windows-mode.
PRGP workload	In MVS/SP 5.0 or earlier, or in compatibility mode in MVS/SP 5.1 or later, composite of service classes. MAINVIEW for OS/390 creates a performance group workload for each performance group defined in the current IEAIPS.xx member.

procedure library	<p>Data set consisting of members that contain executable procedures used by MAINVIEW AutoOPERATOR. These procedures are execute command lists (EXECs) that automate site functions. There can be several versions:</p> <ul style="list-style-type: none"> • the distributed parameter library, called BBPROC • a site-specific parameter library or libraries <p>These can be</p> <ul style="list-style-type: none"> • a library created by AutoCustomization, called UBBPROC • a library created manually, with a unique name <p>The site-created EXECs can be either user-written or customized MAINVIEW AutoOPERATOR-supplied EXECs from BBPROC.</p>
product address space	<p><i>See PAS.</i></p>
profile library	<p>Data set consisting of members that contain profile information and cycle refresh definitions for a terminal session connected to a BBI-SS PAS. Other members are dynamically created by MAINVIEW applications. There can be several versions:</p> <ul style="list-style-type: none"> • the distributed profile library, called BBPROF • a site-specific profile library or libraries <p>These can be</p> <ul style="list-style-type: none"> • a library created by AutoCustomization, called SBBPROF • a library created manually, with a unique name <p>The site library is a common profile shared by all site users. The terminal session CLIST creates a user profile automatically if one does not exist; it is called userid.BBPROF, where userid is your logon ID. User profile libraries allow each user to specify unique PF keys, CYCLE commands, target system defaults, a Primary Option Menu, and a unique set of application profiles.</p>
query	<p>One of two constituent parts of a view; the other is form. A query defines the data for a view; a form defines the display format. <i>See also</i> form, view.</p>
real-time data	<p>Performance data as it exists at the moment of inquiry. Real-time data is recorded during the smallest unit of time for data collection. <i>Contrast with</i> historical data. <i>See also</i> current data and interval data.</p>
Resource Analyzer	<p>Online real-time displays used to analyze IMS resources and determine which are affected by specific workload problems.</p>

Resource Monitor	Online data collection services used to monitor IMS resources and issue warnings when defined utilization thresholds are exceeded.
row	(1) Horizontal component of a view or display comprising all the fields pertaining to a single device, address space, user, and so on. (2) Horizontal component of a DB2 table consisting of a sequence of values, one for each column of the table.
RxD2	Product that provides access to DB2 from REXX. It provides tools to query the DB2 catalog, issue dynamic SQL, test DB2 applications, analyze EXPLAIN data, generate DDL or DB2 utility JCL, edit DB2 table spaces, perform security administration, and much more.
sample cycle	<p>Time between data samples.</p> <p>For the CMF MONITOR Extractor, this is the time specified in the extractor control statements (usually 1 to 5 seconds).</p> <p>For real-time data, the cycle is not fixed. Data is sampled each time you press Enter.</p>
sample library	<p>Data set consisting of members each of which contains one of the following items:</p> <ul style="list-style-type: none"> • sample JCL that can be edited to perform specific functions • macro that is referenced in the assembly of user-written services • sample user exit routine <p>There can be several versions:</p> <ul style="list-style-type: none"> • the distributed sample library, called BBSAMP • a site-specific sample library or libraries <p>These can be</p> <ul style="list-style-type: none"> • a library created by AutoCustomization, called UBBSAMP • a library created manually, with a unique name
sampler	Program that monitors a specific aspect of system performance. Includes utilization thresholds used by the Exception Monitor. The CMF MONITOR Extractor contains samplers.
SBBPROF	<i>See</i> profile library.
scope	Subset of an SSI context. The scope could be all the data for the context or a subset of data within the context. It is user- or site-defined. <i>See</i> SSI context, target.

screen definition	Configuration of one or more views that have been stored with the SAVEScr command and assigned a unique name. A screen includes the layout of the windows and the view, context, system, and product active in each window.
selection view	In MAINVIEW products, view displaying a list of available views.
service class workload	<p>Collection of address spaces defined to OS/390 or z/OS. If you are running Workload Manager (WLM) in goal mode, MAINVIEW for OS/390 creates a service class workload for each service class that you define through WLM definition dialogs.</p> <p>If you are running MVS 4.3 or earlier, or MVS/SP 5.1 or later with WLM in compatibility mode, OS/390 creates a performance group workload instead of a service class. <i>See</i> performance group workload.</p>
service objective	Workload performance goal, specified in terms of response time for TSO workloads or turnaround time for batch workloads. Performance group workloads can be measured by either objective. Composite workload service objectives consist of user-defined weighting factors assigned to each constituent workload. For compatibility mode, neither OS/390 nor z/OS provides any way to measure service.
service point	<p>Specification, to MAINVIEW, of the services required to enable a specific product. Services can be actions, selectors, or views. Each target (for example, CICS, DB2, or IMS) has its own service point.</p> <p>The PLEX view lists all the defined service points known to the CAS to which the terminal session is connected.</p>
service request block (SRB)	Control block that represents a routine to be dispatched. SRB mode routines generally perform work for the operating system at a high priority. An SRB is similar to a task control block (TCB) in that it identifies a unit of work to the system. <i>See also</i> task control block.
service select code	Code entered to invoke analyzers, monitors, and general services. This code is also the name of the individual service.
session	Total period of time an address space has been active. A session begins when monitoring can be performed. If the product address space (PAS) starts after the job, the session starts with the PAS.
SG-Auto	<i>See</i> MAINVIEW SRM SG-Auto.
SG-Control	<i>See</i> MAINVIEW SRM SG-Control.

single system image (SSI)

Feature of the MAINVIEW window environment architecture where you can view and perform actions on multiple OS/390 or z/OS systems as though they were a single system. The rows of a single tabular view can contain rows from different OS/390 or z/OS images.

Skeleton Tailoring Facility

A facility in MAINVIEW AutoOPERATOR that allows skeleton JCL to be used during job submission. Skeleton JCL can contain variables within the JCL statements to be substituted with data values at job submission time. Directive statements can be used in the skeleton JCL to cause the repetition of a set of skeleton statements. This facility functions similar to the TSO skeleton tailoring facility.

SRB

See service request block.

SSI

See single system image.

SSI context

Name created to represent one or more targets for a given product. *See* context, target.

started task workload

Address spaces running jobs that were initiated programmatically.

statistics interval

For MAINVIEW for DB2, cumulative count within a predefined interval (30-minute default set by the DB2STATS parameter in the distributed BBPARM member BBIISP00) for an analyzer service DELTA or RATE display. Specifying the DELTA parameter displays the current value as the difference between the value sampled by the current analyzer request and the value sampled at the start of the current interval. Specifying the RATE parameter displays the current value by minute (DELTA divided by the number of elapsed minutes).

stem variables

A REXX facility, supported in MAINVIEW AutoOPERATOR REXX EXECs and the Skeleton Tailoring Facility, where variable names end with a period followed by a number, such as &POOL.1. This configuration allows each variable to actually represent a table or array of data, with the zero variable containing the number of entries in the array. For example, &POOL.0 = 5 would indicate variables &POOL.1 through &POOL.5 exist.

StopX37/II

See MAINVIEW SRM StopX37/II.

StorageGUARD

See MAINVIEW SRM StorageGUARD.

summary view

View created from a tabular view using the Summarize option in view customization. A summary view compresses several rows of data into a single row based on the summarize criteria.

SYSPROG services	Component of MAINVIEW for OS/390. Over 100 services that detect, diagnose, and correct OS/390 or z/OS system problems as they occur. Accessible from the OS/390 Performance and Control Main Menu. Note that this component is also available as a stand-alone product MAINVIEW SYSPROG Services.
system resource	<i>See</i> object.
target	Entity monitored by one or more MAINVIEW products, such as an OS/390 or z/OS image, an IMS or DB2 subsystem, a CICS region, or related workloads across systems. <i>See</i> context, scope, SSI context.
target context	Single target/product combination. <i>See</i> context.
TASCOSTR	MAINVIEW for IMS Offline program that summarizes detail and summary IMS Resource Utilization Files (IRUFs) to be used as input to the offline components.
task control block (TCB)	Address space-specific control block that represents a unit of work that is dispatched in the address space in which it was created. <i>See also</i> service request block.
TCB	<i>See</i> task control block.
terminal session (TS)	Single point of control for MAINVIEW products, allowing data manipulation and data display and providing other terminal user services for MAINVIEW products. The terminal session runs in a user address space (either a TSO address space or a stand-alone address space for EXCP/VTAM access).
TDIR	<i>See</i> trace log directory.
threshold	Specified value used to determine whether the data in a field meets specific criteria.
TLDS	<i>See</i> trace log data set.
total mode	Usage mode in CMFMON wherein certain columns of data reflect the cumulative value between collection intervals. Invoked by the DELta OFF command. <i>See also</i> collection interval, delta mode.
trace	(1) Record of a series of events chronologically listed as they occur. (2) Online data collection and display services that track transaction activity through DB2, IMS, or CICS.

trace log data set (TLDS)

Single or multiple external VSAM data sets containing summary or detail trace data for later viewing or printing. The trace log(s) can be defined as needed or dynamically allocated by the BBI-SS PAS. Each trace request is assigned its own trace log data set(s).

trace log directory (TDIR)

VSAM linear data set containing one entry for each trace log data set. Each entry indicates the date and time of data set creation, the current status of the data set, the trace target, and other related information.

transaction

Specific set of input data that initiates a predefined process or job.

Transaction Accountant

MVIMS Offline component that produces cost accounting and user charge-back records and reports.

TS

See terminal session.

TSO workload

Workload that consists of address spaces running TSO sessions.

UAS

See user address space.

UBBPARM

See parameter library.

UBBPROC

See procedure library.

UBBSAMP

See sample library.

user address space

Runs a MAINVIEW terminal session (TS) in TSO, VTAM, or EXCP mode.

User BBPROF

See profile library.

view

Formatted data within a MAINVIEW window, acquired from a product as a result of a view command or action. A view consists of two parts: query and form. *See also* form, job activity view, query.

view definition

Meaning of data that appears online, including source of data, selection criteria for data field inclusion and placement, data format, summarization, context, product, view name, hyperlink fields, and threshold conditions.

view command

Name of a view that you type on the COMMAND line to display that view.

view command stack

Internal stack of up to 10 queries. For each command, the stack contains the filter parameters, sort order, context, product, and time frame that accompany the view.

view help	Online help describing the purpose of a view. To display view help, place the cursor on the view name on the window information line and press PF1 (HELP).
window	Area of the MAINVIEW screen in which views and resources are presented. A window has visible boundaries and can be smaller than or equal in size to the MAINVIEW window area. <i>See</i> active window, alternate window, current window, MAINVIEW window area.
window information line	Top border of a window. Shows the window identifier, the name of the view displayed in the window, the system, the scope, the product reflected by the window, and the tomfooleries for which the data in the window is relevant. <i>See also</i> window status field.
window number	Sequential number assigned by MAINVIEW to each window when it is opened. The window number is the second character in the window status field. <i>See also</i> window status field.
window status	One-character letter in the window status field that indicates when a window is ready to receive commands, is busy processing commands, is not to be updated, or contains no data. It also indicates when an error has occurred in a window. The window status is the first character in the window status field. <i>See also</i> window information line, window status field.
window status field	Field on the window information line that shows the current status and assigned number of the window. <i>See also</i> window number, window status.
windows mode	Display of one or more MAINVIEW product views on a screen that can be divided into a maximum of 20 windows. A window information line defines the top border of each window. <i>Contrast with</i> full-screen mode.
WLM workload	In goal mode in MVS/SP 5.1 and later, a composite of service classes. MAINVIEW for OS/390 creates a workload for each WLM workload defined in the active service policy.
workflow	Measure of system activity that indicates how efficiently system resources are serving the jobs in a workload.
workload	(1) Systematic grouping of units of work (for example, address spaces, CICS transactions, IMS transactions) according to classification criteria established by a system administrator. (2) In OS/390 or z/OS, a group of service classes within a service definition.
workload activity view	Tracks workload activity as the workload accesses system resources. A workload activity view measures workload activity in terms of resource consumption and how well the workload activity meets its service objectives.

-
- Workload Analyzer** Online data collection and display services used to analyze IMS workloads and determine problem causes.
- workload definition** Workload created through the WKLIST view. Contains a unique name, a description, an initial status, a current status, and selection criteria by which address spaces are selected for inclusion in the workload. *See* Workload Definition Facility.
- Workload Definition Facility**
In MAINVIEW for OS/390, WKLIST view and its associated dialogs through which workloads are defined and service objectives set.
- workload delay view**
Tracks workload performance as the workload accesses system resources. A workload delay view measures any delay a workload experiences as it contends for those resources.
- Workload Monitor** Online data collection services used to monitor IMS workloads and issue warnings when defined thresholds are exceeded.
- workload objectives**
Performance goals for a workload, defined in WKLIST. Objectives can include measures of performance such as response times and batch turnaround times.

Index

A

AC_CODE 3-60
ACF2USER 3-42
ACS routines
 testing 3-66
affinity separation for STK silos 3-116
ALCTYPE 3-108, 3-109
allocating data sets to DASD volume pools 3-3
ALLTAPE
 intercepting all tape generics 3-32, 3-45,
 3-59, 3-134
assigning tape data sets to pools 3-134
AVL 3-108, 3-109

B

BATCH 3-68
BESTFIT 3-21
BLKINPUT 2-3
BLKOLDSR 2-3
BLKSIZE 3-55
blocksize 3-54
 optimizing for data sets 3-54
BUFSP 3-143, 3-144

C

CAL 3-60
CANDIDATE 3-5, 3-6
CATALOG 3-143, 3-144
changing allocation units to blocks 3-94

CHECK 2-3, 3-4
CISIZE 3-143, 3-144
 statement 3-144
clusters
 VSAM 3-18
COMP 3-129
CONTIG 3-108, 3-110
conventions
 document xiv
 syntax statements xv
 typographical xiv
conversion to blocks 3-94
CRITBIAS 3-5, 3-7, 3-39, 3-40, 3-75, 3-76
CRITDSN 3-7, 3-17, 3-21, 3-84, 3-86
CRITEMC 3-5, 3-7, 3-39, 3-40, 3-75, 3-77
CRITFAIL 3-5, 3-7, 3-39, 3-40, 3-75, 3-77
CRITLIST 2-3, 3-4, 3-5, 3-8, 3-39, 3-41, 3-75,
 3-77

D

DASD performance optimization 3-24, 3-86
DASDGENR 2-3, 3-4
DASDPOOL 3-3
data sets
 allocating with optimal blocksize 3-54
 allocation to virtual storage 3-140
 allocations, limiting space 3-99
 assign tape to pools 3-134
 DASD volume pooling 3-3
 data compaction for tape cartridge 3-129
 deletion, MODDELET 3-48
 empty, initializing 3-50

- expiration date 3-59
- forced cataloging 3-45
- limiting size 3-99, 3-108
- name checking operation 3-34
- permanent 3-99
- release space at close 3-104
- rename requiring management class change 3-70
- size 3-99, 3-108
- temporary 1-1, 3-35, 3-108
 - VIO allocation of 3-140
- utilities 3-106
- DATACLAS 3-67
- DD 3-67
- ddname 3-68
- defer tape mounts 3-132
- DEFUNIT 3-116, 3-117
- delete function
 - MODDELET 3-48
- device 3-68, 3-94, 3-108, 3-129, 3-132, 3-134
- DFSMS enhancements
 - DSNCHECK, data set name checking 3-31
 - OPENEMPT, initialize empty data sets 3-50
 - OPTBLKSZ, allocate with optimal data set blocksize 3-54
 - SMSACSTE, test ACS routines 3-66
 - TAPEPOOL, non-DASD volume pooling 3-134
- DIR 3-108, 3-110
- DMYUNIT 2-3, 3-4
- document conventions xiv
- documentation
 - related xii
- DP_RENAME 2-3, 3-4
- DPO 3-17, 3-21, 3-24, 3-84, 3-86
 - performance data display 3-26, 3-89
 - resources 3-25, 3-87
 - starting and stopping 3-24, 3-87
 - usage 3-25, 3-87
- DPORDEF 3-5, 3-8, 3-75, 3-78
- DPORMAX 3-5, 3-8, 3-25, 3-75, 3-78, 3-87
- DPORMIN 3-5, 3-9, 3-25, 3-75, 3-79, 3-87
- DPORSEP 3-5, 3-9, 3-25, 3-75, 3-79, 3-88
- DPOWIND 3-5, 3-9, 3-25, 3-76, 3-80, 3-87
- DSN 3-67
- DSNCHECK 3-31
 - operational method 3-34
- DSNn 3-32, 3-67

- DSNTYPE 3-67
- DSORG 3-67
- DSTYPE 3-67

E

- EasyPOOL functions
 - SETEXPDT 3-59
- ENVIR 3-67, 3-72
- ERASE 3-143, 3-145
- esoteric name 3-137
- EVENTID 3-5, 3-32, 3-39, 3-41, 3-45, 3-46, 3-48, 3-50, 3-55, 3-60, 3-61, 3-66, 3-70, 3-76, 3-80, 3-95, 3-100, 3-104, 3-105, 3-108, 3-110, 3-116, 3-117, 3-122, 3-125, 3-129, 3-130, 3-132, 3-134, 3-135, 3-140, 3-143, 3-145
- EXPDT 3-60, 3-61, 3-67
- expiration date for data set 3-59

F

- FDRAISIT 3-39
- FORCE 3-55, 3-56
- FORCECAT 3-45
- forced cataloging
 - data sets 3-45
- forcing catalog disposition for new data sets 3-45
- functions
 - DFSMS enhancements
 - DSNCHECK 3-31
 - OPENEMPT 3-50
 - OPTBLKSZ 3-54
 - SMSACSTE 3-66
 - TAPEPOOL 3-134
 - space management
 - OPENEMPT 3-50
 - OPTBLKSZ 3-54
 - SPACCONV 3-94
 - SPACLIMI 3-99
 - SPACSQTY 3-108
 - TAPEDEFR 3-132
 - VIOALLOC 3-140
 - standards control
 - DASDPOOL 3-3
 - FORCECAT 3-45

SUPJSCAT 3-122

G

generic name 3-68, 3-137

H

HDPODAYS 3-5, 3-13, 3-27, 3-76, 3-80, 3-90
HDPOETIM 3-5, 3-12, 3-27, 3-76, 3-81, 3-90
HDPORDEF 3-5, 3-10, 3-33, 3-76, 3-81
HDPORMAX 3-5, 3-11, 3-27, 3-76, 3-81, 3-90
HDPORMIN 3-5, 3-11, 3-27, 3-76, 3-82, 3-90
HDPORSEP 3-5, 3-12, 3-27, 3-76, 3-82, 3-90
HDPOSTIM 3-5, 3-12, 3-27, 3-76, 3-90
HISTDAYS 2-3, 3-4, 3-27, 3-75, 3-90
HISTDPO 3-17, 3-21, 3-26, 3-27, 3-85, 3-89,
3-90
historical DASD performance optimization 3-26
HLQ 3-67

I

IDRC 3-129, 3-134
IMBED 3-143, 3-145
initializing empty data sets 3-50
ISAM 3-56
ISPF interface
performance data by pool 3-24, 3-87
ISPF service 3.2 3-106

J

JCLEXT 2-3, 3-4
JCLUREQ 2-3, 3-4
JOB 3-67
JOBACCTn 3-67
JOBCAT 3-122
JOBCAT DD
suppressing 3-122

L

LIMIT 3-100

limiting space requested by users 3-99
LLQ 3-32, 3-33, 3-67

M

management class
data set rename 3-70
MAXQLF 3-32, 3-33
MAXSIZE 3-18, 3-67, 3-99
MAXSPACE 3-17, 3-21, 3-85
MGMTCLAS 3-67, 3-72
MINQLF 3-32, 3-34
MODDELET 3-48

N

non-DASD volume pooling 3-134
NQUAL 3-67
NVOL 3-6, 3-13, 3-21, 3-134
NVOLINDX 3-6, 3-13, 3-21
NVOLMAX 3-5, 3-13, 3-21

O

OPENEMPT 3-50
opening empty data sets 3-50
OPTBLKSZ 3-54
optimum blocksize 3-54
OWNER 3-143, 3-145

P

PERCENT 3-17, 3-21, 3-85
permanent data sets 3-99
PGM 3-67
POOL 3-6, 3-14, 3-39, 3-41, 3-134, 3-135
pool
non-DASD volumes 3-134
volume pooling for DASD data sets 3-3
POOLNAME 5-2
description 5-2
PQTY 3-108, 3-111
PROCOLD 2-3, 3-5
PWDDEL 3-143, 3-146

R

RACF 3-67, 3-68
RAIDDEVTYPE 3-6, 3-14, 3-76, 3-82
RECORD 3-67
REJECT 2-3, 3-5, 3-6, 3-15, 3-20, 3-134, 3-136
related publications xii
release notes xiii
releasing data set space at close 3-104
REPL 3-143, 3-146
REPLACE 3-60, 3-62, 3-108, 3-111
response time 3-24, 3-87
RETPD 3-60, 3-62, 3-67
REUSE 3-143
RLSE 3-104, 3-105
ROUND 3-95, 3-108, 3-111

S

selecting volume from storage group 3-74
SEP 3-6, 3-15
SETEXPDT 3-59
setting data compaction for tape cartridge data set 3-129
setting expiration date for data set 3-59
setting optimum blocksize 3-54
setting primary and secondary space for data set 3-108
SIZE 3-67
SMPOOLxx subparameters
 POOLNAME 5-2
SMS_ALLOC 2-3, 3-75
SMS_EXTEND 2-3, 3-75
SMSACSTE 3-66
SMSMANAGED 3-66, 3-67
SMSMCREN 3-70
SMSPOOL 3-76, 3-83
SMSPOOL_EXT 3-76, 3-84
SMSSELECT 3-74
SPACCONV 3-74, 3-94
space allocation requests 3-99
space management
 OPENEMPT, initialize empty data sets 3-50
 OPTBLKSZ, allocate with optimal blocksize 3-54
 SPACLIMI, limits space allocation requests 3-99

SPACSQTY, set primary and secondary space for data sets 3-108
TAPEDEFR, assign defer to tape allocations 3-132
VIOALLAC, control VIO allocation 3-140
SPACLIMI 3-99
SPACRLSE 3-104
SPACSQTY 3-108
SPACVOLA 3-6, 3-20, 3-125
SPLIT 3-116, 3-117
SQTY 3-108, 3-112
standards control
 DASDPOOL, volume pooling for DASD data sets 3-3
 FORCECAT, forced cataloging 3-45
 SUPJSCAT, suppressing JOBCAT, SUBCAT DDs 3-122
 TAPEPOOL, non-DASD volume pooling 3-134
STC 3-68
STEPACCTn 3-67
STEPCAT 3-122
STKSCR 2-3
STKSUPP 3-116
STOGROUP 3-67
STORCLAS 3-67
SUPJSCAT 3-122
suppressing JOBCAT and STEPCAT DDs 3-122
suppressing volume references 3-125
SUPVOL 3-6, 3-15, 3-23
SUPVOLRF 3-125
syntax statement conventions xv

T

TAPECOMP 3-129
TAPEDEFR 3-132
TAPEGENR 2-3
TAPEPOOL 3-134
temporary data sets 1-1, 3-35, 3-108, 3-140
test ACS routines 3-66
trace 3-67
TRKCYL 2-3, 3-95, 3-96
TRKLEN 2-3, 3-95, 3-96
TSO 3-47, 3-68, 3-103
typographical conventions xiv

U

UNIT 3-6, 3-15, 3-18, 3-21, 3-67, 3-135, 3-136
USELIMIT 3-22
USER 3-67
USEVOL 3-6, 3-16, 3-20, 3-39, 3-42, 3-135,
3-136

V

VCOMPLLQ 3-32
 defined 3-34
VFORCE 3-6, 3-16, 3-21
VIO 3-140, 3-141
 allocating of temporary data sets 3-140
 controlling allocation 3-140
VIOALLOC 3-140
VOL 3-67, 5-2
VOLSEL 3-6, 3-16, 3-21, 3-27, 3-39, 3-42,
3-76, 3-84, 3-90
volume pooling
 non-DASD 3-134
volume selection 3-6, 3-16, 3-21, 3-42, 3-54,
3-55, 3-84, 3-94
VSAM 3-5, 3-18, 3-34, 3-51, 3-56, 3-62, 3-96,
3-101
VSAMCNTL 3-143

X

XMODE 3-67

STOP!

IMPORTANT INFORMATION - DO NOT INSTALL THIS PRODUCT UNLESS YOU HAVE READ ALL OF THE FOLLOWING MATERIAL

By clicking the YES or ACCEPT button below (when applicable), or by installing and using this Product or by having it installed and used on your behalf, You are taking affirmative action to signify that You are entering into a legal agreement and are agreeing to be bound by its terms, EVEN WITHOUT YOUR SIGNATURE. BMC is willing to license this Product to You ONLY if You are willing to accept all of these terms. CAREFULLY READ THIS AGREEMENT. If You DO NOT AGREE with its terms, DO NOT install or use this Product; press the NO or REJECT button below (when applicable) or promptly contact BMC or your BMC reseller and your money will be refunded if by such time You have already purchased a full-use License.

SOFTWARE LICENSE AGREEMENT FOR BMC PRODUCTS

SCOPE. This is a legally binding Software License Agreement ("**License**") between You (either an individual or an entity) and BMC pertaining to the original computer files (including all computer programs and data stored in such files) contained in the enclosed Media (as defined below) or made accessible to You for electronic delivery, if as a prerequisite to such accessibility You are required to indicate your acceptance of the terms of this License, and all whole or partial copies thereof, including modified copies and portions merged into other programs (collectively, the "**Software**"). "**Documentation**" means the related hard-copy or electronically reproducible technical documents furnished in association with the Software, "**Media**" means the original BMC-supplied physical materials (if any) containing the Software and/or Documentation, "**Product**" means collectively the Media, Software, and Documentation, and all Product updates subsequently provided to You, and "**You**" means the owner or lessee of the hardware on which the Software is installed and/or used. "**BMC**" means BMC Software Distribution, Inc. unless You are located in one of the following regions, in which case "BMC" refers to the following indicated BMC Software, Inc. subsidiary: (i) Europe, Middle East or Africa --BMC Software Distribution, B.V., (ii) Asia/Pacific -- BMC Software Asia Pacific Pte Ltd., (iii) Brazil -- BMC Software do Brazil, or (iv) Japan -- BMC Software K.K. If You enter into a separate, written software license agreement signed by both You and BMC or your authorized BMC reseller granting to you the rights to install and use this Product, then the terms of that separate, signed agreement will apply and this License is void.

FULL-USE LICENSE. Subject to these terms and payment of the applicable license fees, BMC grants You this non-exclusive License to install and use one copy of the Software for your internal use on the number(s) and type(s) of servers or workstations for which You have paid or agreed to pay to BMC or your BMC reseller the appropriate license fee. If your license fee entitles You only to a License having a limited term, then the duration of this License is limited to that term; otherwise this License is perpetual, subject to the termination provisions below.

TRIAL LICENSE. If You have not paid or agreed to pay to BMC or your BMC Reseller the appropriate license fees for a full use license, then, **NOTWITHSTANDING ANYTHING TO THE CONTRARY CONTAINED IN THIS LICENSE:** (i) this License consists of a non-exclusive evaluation license ("Trial License") to use the Product for a limited time ("Trial Period") only for evaluation; (ii) during the Trial Period, You may not use the Software for development, commercial, production, database management or other purposes than those expressly permitted in clause (i) immediately above; and (iii) your use of the Product is on an **AS IS** basis, and **BMC, ITS RESELLERS AND LICENSORS GRANT NO WARRANTIES OR CONDITIONS (INCLUDING IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE) TO YOU AND ACCEPT NO LIABILITY WHATSOEVER RESULTING FROM THE USE OF THIS PRODUCT UNDER THIS TRIAL LICENSE.** If You use this Product for other than evaluation purposes or wish to continue using it after the Trial Period, you must purchase a full-use license. When the Trial Period ends, your right to use this Product automatically expires, though in certain cases You may be able to extend the term of the Trial Period by request. Contact BMC or your BMC reseller for details.

TERM AND TERMINATION. This License takes effect on the first to occur of the date of shipment or accessibility to You for electronic delivery, as applicable (the "**Product Effective Date**"). You may terminate this License at any time for any reason by written notice to BMC or your BMC reseller. This License and your right to use the Product will terminate automatically with or without notice by BMC if You fail to comply with any material term of this License. Upon termination, You must erase or destroy all components of the Product including all copies of the Software, and stop using or accessing the Software. Provisions concerning Title and Copyright, Restrictions (or Restricted Rights, if You are a U.S. Government entity) or limiting BMC's liability or responsibility shall survive any such termination.

TITLE AND COPYRIGHT; RESTRICTIONS. All title and copyrights in and to the Product, including but not limited to all modifications thereto, are owned by BMC and/or its affiliates and licensors, and are protected by both United States copyright law and applicable international copyright treaties. You will not claim or assert title to or ownership of the Product. To the extent expressly permitted by applicable law or treaty notwithstanding this limitation, You may copy the Software only for backup or archival purposes, or as an essential step in utilizing the Software, but for no other purpose. You will not remove or alter any copyright or proprietary notice from copies of the Product. You acknowledge that the Product contains valuable trade secrets of BMC and/or its affiliates and licensors. Except in accordance with the terms of this License, You agree (a) not to decompile, disassemble, reverse engineer or otherwise attempt to derive the Software's source code from object code except to the extent expressly permitted by applicable law or treaty despite this limitation; (b) not to sell, rent, lease, license, sublicense, display, modify, time share, outsource or otherwise transfer the Product to, or permit the use of this Product by, any third party; and (c) to use reasonable care and protection to prevent the unauthorized use, copying, publication or dissemination of the Product and BMC confidential information learned from your use of the Product. **You will not export or re-export any Product without both the written consent of BMC and the appropriate U.S. and/or foreign government license(s) or license exception(s).** Any programs, utilities, modules or other software or documentation created, developed, modified or enhanced by or for You using this Product shall likewise be subject to these restrictions. BMC has the right to obtain injunctive relief against any actual or threatened violation of these restrictions, in addition to any other available remedies. Additional restrictions may apply to certain files, programs or data supplied by third parties and embedded in the Product; consult the Product installation instructions or Release Notes for details.

LIMITED WARRANTY AND CONDITION. If You have purchased a Full-Use License, BMC warrants that (i) the Media will be, under normal use, free from physical defects, and (ii) for a period of ninety (90) days from the Product Effective Date, the Product will perform in substantial accordance with the operating specifications contained in the Documentation that is most current at the Product Effective Date. BMC's entire liability and your exclusive remedy under this provision will be for BMC to use reasonable best efforts to remedy defects covered by this warranty

and condition within a reasonable period of time or, at BMC's option, either to replace the defective Product or to refund the amount paid by You to license the use of the Product. BMC and its suppliers do not warrant that the Product will satisfy your requirements, that the operation of the Product will be uninterrupted or error free, or that all software defects can be corrected. This warranty and condition shall not apply if: (i) the Product is not used in accordance with BMC's instructions, (ii) a Product defect has been caused by any of your or a third party's malfunctioning equipment, (iii) any other cause within your control causes the Product to malfunction, or (iv) You have made modifications to the Product not expressly authorized in writing by BMC. No employee, agent or representative of BMC has authority to bind BMC to any oral representations, warranties or conditions concerning the Product. **THIS WARRANTY AND CONDITION IS IN LIEU OF ALL OTHER WARRANTIES AND CONDITIONS. THERE ARE NO OTHER EXPRESS OR IMPLIED WARRANTIES OR CONDITIONS, INCLUDING THOSE OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE, REGARDING THIS LICENSE OR ANY PRODUCT LICENSED HEREUNDER. THIS PARAGRAPH SHALL NOT APPLY TO A TRIAL LICENSE.** Additional support and maintenance may be available for an additional charge; contact BMC or your BMC reseller for details.

LIMITATION OF LIABILITY. Except as stated in the next succeeding paragraph, BMC's and your BMC reseller's total liability for all damages in connection with this License is limited to the price paid for the License. **IN NO EVENT SHALL BMC BE LIABLE FOR ANY CONSEQUENTIAL, SPECIAL, INCIDENTAL, PUNITIVE OR INDIRECT DAMAGES OF ANY KIND ARISING OUT OF THE USE OF THIS PRODUCT (SUCH AS LOSS OF PROFITS, GOODWILL, BUSINESS, DATA OR COMPUTER TIME, OR THE COSTS OF RECREATING LOST DATA), EVEN IF BMC HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.** Some jurisdictions do not permit the limitation of consequential damages so the above limitation may not apply.

INDEMNIFICATION FOR INFRINGEMENT. BMC will defend or settle, at its own expense, any claim against You by a third party asserting that your use of the Product within the scope of this License violates such third party's patent, copyright, trademark, trade secret or other proprietary rights, and will indemnify You against any damages finally awarded against You arising out of such claim. However, You must promptly notify BMC in writing after first receiving notice of any such claim, and BMC will have sole control of the defense of any action and all negotiations for its settlement or compromise, with your reasonable assistance. BMC will not be liable for any costs or expenditures incurred by You without BMC's prior written consent. If an order is obtained against your use of the Product by reason of any claimed infringement, or if in BMC's opinion the Product is likely to become the subject of such a claim, BMC will at its option and expense either (i) procure for You the right to continue using the product, or (ii) modify or replace the Product with a compatible, functionally equivalent, non-infringing Product, or (iii) if neither (i) nor (ii) is practicable, issue to You a pro-rata refund of your paid license fee(s) proportionate to the number of months remaining in the 36 month period following the Product Effective Date. This paragraph sets forth your only remedies and the total liability to You of BMC, its resellers and licensors arising out of such claims.

GENERAL. This License is the entire understanding between You and BMC concerning this License and may be modified only in a mutually signed writing between You and BMC. If any part of it is invalid or unenforceable, that part will be construed, limited, modified, or severed so as to eliminate its invalidity or unenforceability. This License will be governed by and interpreted under the laws of the jurisdiction named below, without regard to conflicts of law principles, depending on which BMC Software, Inc. subsidiary is the party to this License: (i) BMC Software Distribution, Inc. - the State of Texas, U.S.A., (ii) BMC Software Distribution, B.V. - The Netherlands, (iii) BMC Software Asia Pacific Pte Ltd. -- Singapore (iv) BMC Software do Brazil -- Brazil, or (v) BMC Software K.K. -- Japan. Any person who accepts or signs changes to the terms of this License promises that they have read and understood these terms, that they have the authority to accept on your behalf and legally obligate You to this License. Under local law and treaties, the restrictions and limitations of this License may not apply to You; You may have other rights and remedies, and be subject to other restrictions and limitations.

U.S. GOVERNMENT RESTRICTED RIGHTS. UNPUBLISHED -- RIGHTS RESERVED UNDER THE COPYRIGHT LAWS OF THE UNITED STATES. Use, duplication, or disclosure by the U.S. Government is subject to restrictions set forth in FAR Section 52.227-14 Alt. III (g)(3), FAR Section 52.227-19, DFARS 252.227-7014 (b) or DFARS 227.7202, as amended from time to time. Contractor/Manufacturer is BMC Software, Inc., 2101 CityWest Blvd., Houston, TX 77042-2827, USA. Any contract notices should be sent to this address.

Notes



100042496